

NRCL - A Model Building Approach to the Bernays-Schönfinkel Fragment (Full Paper)

Gábor Alagi^{1,2,3} and Christoph Weidenbach³

¹ Saarbrücken Graduate School of Computer Science, Germany

² Saarland University, Saarbrücken, Germany

³ Max-Planck-Institut für Informatik, Saarbrücken, Germany
`{galagi, weidenbach}@mpi-inf.mpg.de`

Abstract

We combine constrained literals for model representation with key concepts from first-order superposition and propositional conflict-driven clause learning (CDCL) to create the new calculus *Non-Redundant Clause Learning* (NRCL) deciding the Bernays-Schönfinkel fragment. Our calculus uses first-order literals constrained by disequations between tuples of terms for compact model representation. From superposition, NRCL inherits the abstract redundancy criterion and the monotone model operator. CDCL adds the dynamic, conflict-driven search for an atom ordering inducing a model. As a result, in NRCL a false clause can be found effectively modulo the current model candidate. It guides the derivation of a first-order ordered resolvent that is never redundant. Similar to 1UIP-learning in CDCL, the learned resolvent induces backtracking and, by blocking the previous conflict state via propagation, it enforces progress towards finding a model or a refutation. The non-redundancy result also implies that only finitely many clauses can be generated by NRCL on the Bernays-Schönfinkel fragment, which serves as an argument for termination.

1 Introduction

The Bernays-Schönfinkel fragment, also called Effectively Propositional Logic, or BS (or EPR) for short, is an important fragment of classic first-order logic, where only constants are allowed as function symbols in the clause normal form.

This decidable and NEXPTIME-complete fragment has many applications, including knowledge representation [19] and ontological reasoning [38], hardware verification [20][30][13], logic programming [12], and planning [32].

Over the years a number of calculi have attempted to provide an efficient solution for BS problems. These approaches range from the early *SEM* and *Mace* systems [39] to the recent state-of-the-art solvers like *iProver* [16] and *Darwin* [7], but even general purpose first-order theorem provers provide specialized techniques for BS problems, like *generalisation* in *Vampire* [31], or specialized splitting techniques for *SPASS* introduced in [18] and [15].

In this paper, we introduce a new calculus for solving BS problems with iterative model building. Our approach builds first-order candidate models instead of approximations, uses constrained literals for model representation, and learns new non-redundant clauses to guide the search.

Our calculus, called NRCL or *Non-Redundant Clause Learning*, shares many principles with propositional SAT solving and superposition. For a detailed introduction to *conflict-driven clause learning (CDCL)*, see the early article [36], or the more recent handbook [8]. The interested reader can get a thorough overview of superposition in the articles [3][4][41][2][28].

Compared to the existing approaches, we use a more expressive and implicit constraint language, our search is guided by backjumping and learning non-redundant clauses, and our model representation is more compact, in general. In addition, compared to all existing approaches, we can prove that all our learned clauses are non-redundant and this way, for the first time, establish a calculus that combines the search with respect to a dynamically changing (partial) model with an overall notion of redundancy. For a more detailed comparison, see Section 9.

In the rest of the paper, we first introduce some basic definitions and notions in Section 2, followed by a description of our calculus in Section 3. Section 4 establishes its soundness, while, after introducing some regularity conditions in Section 5, we provide our key result, namely non-redundant clause learning, in Section 6. We then prove termination in Section 7.

In Section 8, we specify some details on handling constraints, and basic heuristics for a future implementation. We compare our calculus to the existing literature in more details in Section 9. Finally, Section 10 provides a summary and outlines future work.

2 Preliminaries

2.1 Basic Definitions

We assume the reader is familiar with first-order logic, its syntax, and its semantics. In particular, we handle the Bernays-Schönfinkel fragment, or BS for short. In this fragment the only functions allowed in the clause normal form are finitely many constants. We denote the finite *signature* by Σ , the *set of predicate symbols* by Pr , and call the finite set of constants the *domain*, denoted by \mathcal{D} .

We denote the set of all first-order atoms over a signature Σ and a possibly infinite set of variables X by $\mathcal{A}_\Sigma(X)$. In particular the set of ground atoms is denoted by \mathcal{A}_Σ , a short-hand for $\mathcal{A}_\Sigma(\emptyset)$. For a literal L , $|L|$ denotes the atom contained by L . In general, we denote the ground instances of an *expression* - a term, literal, or clause - e over the domain \mathcal{D} by the notation $\text{gnd}(e)$.

W.l.o.g., we assume that each independent expression is variable disjoint, and we call a variable *fresh* if it does not occur in any expression - e.g. clause or clause set - of the current context.

We consider *substitutions* in the usual way, and for a substitution σ , $\text{dom}(\sigma)$ denotes the *domain* of σ , i.e. the finite set of variables with $x \neq x\sigma$, and $\text{rng}(\sigma)$ denotes the *range* of σ , i.e. the image of $\text{dom}(\sigma)$ w.r.t. σ .

We assume the reader is familiar with *most general unifiers*, and mgu is used to denote the result of unifying two or more expressions or substitutions.

We use the short-hand $\exists\sigma = \text{mgu}(e_1, e_2)$ to both state the existence of a most general unifier and bind σ to one such substitution.

For expressions or substitutions e_1, e_2 , we say e_2 *can be matched against* e_1 , or e_1 *is more general than* e_2 , and write $e_1 \geq e_2$, if and only if there is a substitution σ such that $e_1 = e_2\sigma$.

We represent a first-order interpretation I with the set $\{A \in \mathcal{A}_\Sigma \mid I \models A\}$. We define *satisfiability* and *semantic consequence* as usual.

In particular, we consider the problem of deciding whether a finite clause set N over a BS language Σ without equality is satisfiable. This problem is known to be NEXPTIME-complete [25].

2.2 Constraints and Constrained Literals

Next, we provide details about the constraint language we use. Our constraints are equivalent with *implicit generalizations*, a constraint language for representing terms and models with exceptions. It has applications in inductive learning, logic programming and term rewriting. For more details see e.g. [11][24].

The name *dismatching constraints* was chosen in the spirit of *iProver*[22], although for our purposes checking satisfiability has to be carried out over the ground instances and thus, the linear-time algorithm of *iProver* based on matching is not applicable.

While implicit generalizations maintain a list of literals with fresh variables representing exceptions for the literal constrained, dismatching constraints extract the arguments of the literals and represent the restrictions as conjunctions of disequations to allow more simplification and a more compact representation. In particular, we maintain a strict normal form, which already assumes most inexpensive simplifications.

We chose dismatching constraints for a balance between expressiveness and simplicity, for the existing literature, and for compactness. However, NRCL is compatible with any constraint language allowing the operations discussed in the next subsection.

Definition 2.1 (Dismatching Constraint) A dismatching constraint π is of the form

$$\bigwedge_{i \in \mathcal{I}} \vec{s}_i \neq \vec{t}_i$$

where \mathcal{I} is a finite set of indices, and for each $i \in \mathcal{I}$, \vec{s}_i and \vec{t}_i are tuples of terms of the same dimension.

Furthermore, we assume that all the left-hand side variables in π differ from any right-hand side variable, and for each $i, j \in \mathcal{I}$, \vec{t}_i and \vec{t}_j are variable disjoint whenever i differs from j .

We further extend the set of constraints with the constants \top, \perp representing the tautological and the unsatisfiable constraint, respectively.

Finally, an atomic constraint $\vec{s} \neq \vec{t}$ occurring in π is also called a subconstraint of π .

To enforce a normal form, we make further assumptions below.

Definition 2.2 (Normal Form) We say a constraint $\pi = \bigwedge_{i \in \mathcal{I}} \vec{s}_i \neq \vec{t}_i$ is in normal form iff the following conditions hold:

(C1) each \vec{s}_i contains only variables

(C2) no variable occurs more than once in any left-hand side \vec{s}_i

A simple consequence of the normal form is that the two sides of any sub-constraint $\vec{s} \neq \vec{t}$ are always unifiable, and the *induced substitution* $\{\vec{s} \leftarrow \vec{t}\}$ is always well-defined and matches the left-hand side against the right-hand side.

Definition 2.3 (Induced Substitutions) *The set of induced substitutions of a dismatching constraint π in normal form is the set given by*

$$\{\{\vec{s}_i \leftarrow \vec{t}_i\} \mid i \in \mathcal{I}\}$$

if $\pi = \bigwedge_{i \in \mathcal{I}} \vec{s}_i \neq \vec{t}_i$. For \perp , we define it as the set containing only the identity, and for \top as the empty set.

We define $\text{lvar}(\pi)$ and $\text{rvar}(\pi)$ as the set of the left-hand side and right-hand side variables of some dismatching constraint π , respectively. Then the semantics for our constraints can be given as below.

Definition 2.4 *A solution of a constraint π over some variable set V , which contains $\text{lvar}(\pi)$ but contains no variable from $\text{rvar}(\pi)$, is a ground substitution $\delta : V \rightarrow \mathcal{D}$ such that no \vec{t}_i can be matched against the respective $\vec{s}_i\delta$, i.e. no $\vec{s}_i\delta$ is an instance of the respective \vec{t}_i .*

In particular, if $\pi = \top$, any such grounding substitution is a solution, and $\pi = \perp$ has no solution at all.

As usual, π is called *satisfiable* and *unsatisfiable* if it has a solution or no solution, respectively. We note that the notion of satisfiability depends only on $\text{lvar}(\pi)$.

Example 2.5 *Consider the domain $\mathcal{D} = \{a, b\}$ and the constraint*

$$\pi = (x, y) \neq (v, v) \wedge y \neq a$$

Then π is satisfiable and the ground substitution $\sigma = \{x \leftarrow a, y \leftarrow b\}$ is the only solution of π (over $V = \{x, y\}$), since y can only be b and the first subconstraint represents $x \neq y$.

Remark 2.6 *It can be shown that a ground substitution $\delta : V \rightarrow \mathcal{D}$ with $\text{lvar}(\pi) \subseteq V$ is not a solution of π if and only if there is an induced substitution σ which is more general than δ .*

Definition 2.7 *Let π and π' denote constraints for which both*

- $\text{lvar}(\pi) \cap \text{rvar}(\pi') = \emptyset$, and
- $\text{lvar}(\pi') \cap \text{rvar}(\pi) = \emptyset$

hold. Such constraints are called equivalent iff their sets of solutions coincide for any V such that $\text{lvar}(\pi) \cup \text{lvar}(\pi') \subseteq V$, and both $V \cap \text{rvar}(\pi) = \emptyset$ and $V \cap \text{rvar}(\pi') = \emptyset$.

Normal Form Transformation

Next, we show that any dismatching constraint of the form $\bigwedge_{i \in \mathcal{I}} \vec{s}_i \neq \vec{t}_i$ can be normalized in polynomial time. This can be achieved with the rule set below, given as rewriting rules over the subconstraints.

1. $(\vec{s}_1, a, \vec{s}_2) \neq (\vec{t}_1, a, \vec{t}_2) \Rightarrow (\vec{s}_1, \vec{s}_2) \neq (\vec{t}_1, \vec{t}_2)$, where $a \in \mathcal{D}$
2. $(\vec{s}_1, a, \vec{s}_2) \neq (\vec{t}_1, b, \vec{t}_2) \Rightarrow \top$, where $a \neq b \in \mathcal{D}$
3. $(\vec{s}_1, a, \vec{s}_2) \neq (\vec{t}_1, x, \vec{t}_2) \Rightarrow (\vec{s}_1, \vec{s}_2) \neq (\vec{t}_1, \vec{t}_2)\sigma$, where $a \in \mathcal{D}$, $\sigma = \{x \leftarrow a\}$
4. $() \neq () \Rightarrow \perp$
5. $(\vec{s}_1, x, \vec{s}_2, x, \vec{s}_3) \neq (\vec{t}_1, r_1, \vec{t}_2, r_2, \vec{t}_3) \Rightarrow (\vec{s}_1, x, \vec{s}_2, \vec{s}_3) \neq (\vec{t}_1, r_1, \vec{t}_2, \vec{t}_3)\sigma$, if $\exists \sigma = \text{mgu}(r_1, r_2)$
6. $(\vec{s}_1, x, \vec{s}_2, x, \vec{s}_3) \neq (\vec{t}_1, r_1, \vec{t}_2, r_2, \vec{t}_3) \Rightarrow \top$, if $\nexists \text{mgu}(r_1, r_2)$
7. $\vec{s} \neq \vec{t} \Rightarrow \perp$, if \vec{t} can be matched against \vec{s}
8. $(\vec{s}_1, \vec{s}_2) \neq (\vec{t}_1, \vec{t}_2) \Rightarrow \vec{s}_1 \neq \vec{t}_1$, if $\text{var}(\vec{t}_1) \cap \text{var}(\vec{t}_2) = \emptyset$, and \vec{t}_2 can be matched against \vec{s}_2

Where the last rule is considered modulo permutations of positions corresponding to the (\vec{s}_1, \vec{s}_2) -partitionings.

Example 2.8 *Let us normalize the following constraint:*

$$(x, a, y, x) \neq (b, v, w, w) \wedge (x, a, y, x) \neq (w_0, w_0, v_0, t_0)$$

For the first subconstraint we get

$$(x, a, y, x) \neq (b, v, w, w) \stackrel{(3)}{\Rightarrow} (x, y, x) \neq (b, w, w) \stackrel{(4)}{\Rightarrow} (x, y) \neq (b, b)$$

and for the second one

$$(x, a, y, x) \neq (w_0, w_0, v_0, t_0) \stackrel{(4)}{\Rightarrow} (x, a, y) \neq (w_0, w_0, v_0) \stackrel{(3)}{\Rightarrow} (x, y) \neq (a, v_0) \stackrel{(8)}{\Rightarrow} x \neq a$$

Thus, the normalized constraint is

$$(x, y) \neq (b, b) \wedge x \neq a$$

Applying these rules together with the usual rules for conjunction and the constants \perp, \top

1. preserves the variable disjointness conditions of Definition 2.1
2. preserves solutions, i.e. the left-hand side and right-hand side constraints are equivalent
3. transforms π into normal form in polynomial time

We note that the rules (7) and (8) are optional, and that (7) is a special case of (8).

Therefore, w.l.o.g. we assume that the constraints are always in normal form, and the result of any operation is transformed into normal form without explicitly expressing it. We also express it by using the notation $\bigwedge_{i \in \mathcal{I}} \vec{x}_i \neq \vec{t}_i$ for dismatching constraints in the rest of the paper.

Constrained Literals

Next, we define literals constrained with dismatching constraints in normal form, and give their semantics as sets of ground literals.

Definition 2.9 (Constrained Literal) *We call the pair $(L; \pi)$ of a literal L and a dismatching constraint π such that both $\text{lvar}(\pi) \subseteq \text{var}(L)$ and $\text{rvar}(\pi) \cap \text{var}(L) = \emptyset$ hold a constrained literal.*

The semantics of constrained literals is given by the following definition of the set of covered literals:

$$\text{gnd}(L; \pi) = \{L\delta \mid \delta : \text{var}(L) \rightarrow \mathcal{D} \text{ s.t. } \forall i \in \mathcal{I} : \# \text{mgu}(\vec{x}_i\delta, \vec{t}_i)\}$$

where $\pi = \bigwedge_{i \in \mathcal{I}} \vec{x}_i \neq \vec{t}_i$. A ground literal L' is covered by a constrained literal $(L; \pi)$ iff $L' \in \text{gnd}(L; \pi)$.

We say that a constrained literal $(L; \pi)$ is empty if it covers no ground literals, i.e. $\text{gnd}(L; \pi)$ is empty.

It is easy to see that $(L; \pi)$ is empty if and only if π is unsatisfiable, and that given a solution δ of π over $\text{lvar}(\pi)$, for any extension δ' of δ to $\text{var}(L)$, $L\delta' \in \text{gnd}(L; \pi)$ holds.

Example 2.10 Let $(L; \pi) = (P(x, y); (x, y) \neq (v, v) \wedge x \neq a \wedge y \neq b)$. Then the set of covered literals over the domain $\mathcal{D}_2 = \{a, b\}$ is

$$\text{gnd}(L; \pi) = \{P(b, a)\}$$

and if we take $\mathcal{D}_3 = \{a, b, c\}$ instead, it is

$$\text{gnd}(L; \pi) = \{P(b, a), P(c, a), P(b, c)\}$$

In the rest of the paper we make some further assumptions as common in automated reasoning:

1. Different constrained literals are variable disjoint, unless stated otherwise.
2. Apart from normal form transformations, for any substitution σ applied to a constrained literal $(L; \pi)$, the following always hold unless stated otherwise:

- $\text{dom}(\sigma) \cap \text{rvar}(\pi) = \emptyset$
- $\text{var}(\text{rng}(\sigma)) \cap \text{rvar}(\pi) = \emptyset$

Constrained Clauses

Occasionally, we have to represent a collection of ground clauses by a constrained clause $(C; \pi)$. Extending the notations and semantics for constrained literals to constrained clauses is straightforward.

Furthermore, we use the notation $(C; \sigma; \pi)$ for the constrained clause $(C\sigma; \pi)$, whenever we wish to syntactically distinguish C and σ .

We only note that during resolving away literals from C , we might get to a state where $\text{lvar}(\pi)$ contains variables not occurring in C . See the constrained unit clause

$$(P(y, z); (x, y) \neq (v, v) \wedge (x, z) \neq (w, w))$$

from Example 8.1 for a demonstration.

For semantic purposes, these *free variables* are considered existential variables. We assume that such variables are eliminated through instantiation, see Section 8 for further details.

2.3 Operations on Constrained Literals

In the context of our calculus, three operations are of significance: conjunction, difference, and checking whether a constrained literal is empty.

In the literature checking emptiness also relates to *sufficient completeness* and *negation elimination* and it is known to be a co-NP-complete problem [24] in the case of finitely many function symbols and infinite Herbrand universe.

This complexity result also holds for our setting - one might take a binary domain with *true* and *false*, and then each atomic constraint with constant right-hand side can be seen as clauses with the left-hand side variables as propositional variables, and the emptiness of the whole constraint as the unsatisfiability of this clause set.

In this section, we propose an enumeration-based algorithm to test emptiness as an alternative to relying on external CSP and CDCL solvers.

Conjunction

For two constrained literals $(L_1; \pi_1)$, $(L_2; \pi_2)$ with the same polarity and predicate symbol, we look for a constrained literal $(L; \pi)$ for which

$$\text{gnd}(L; \pi) = \text{gnd}(L_1; \pi_1) \cap \text{gnd}(L_2; \pi_2)$$

holds. If the two literals are unifiable, such a literal exists. Otherwise, any empty constrained literal can be chosen.

Definition 2.11 (Conjunction) *Let us define and denote the conjunction of two constrained literals $(L_1; \pi_1)$, $(L_2; \pi_2)$ as*

$$(L_1; \pi_1) \wedge (L_2; \pi_2) = (L_1\sigma; \pi_1\sigma \wedge \pi_2\sigma)$$

if $\exists \sigma = \text{mgu}(L_1, L_2)$. If the literals are not unifiable, we define it as the empty $(L_1; \perp)$.

This definition is sound, i.e.

Lemma 2.12 *For any unifiable constrained literals $(L_1; \pi_1)$, $(L_2; \pi_2)$,*

$$\text{gnd}(L_1; \pi_1) \cap \text{gnd}(L_2; \pi_2) = \text{gnd}(L_1\sigma; \pi_1\sigma \wedge \pi_2\sigma)$$

holds, where $\sigma = \text{mgu}(L_1, L_2)$.

Proof:

(\subseteq): Consider a ground literal from $\text{gnd}(L_1; \pi_1) \cap \text{gnd}(L_2; \pi_2)$, and w.l.o.g. assume it has the form $L_1\delta$. Then $L_1\delta \geq L_i$ holds for both $i = 1, 2$.

Thus, $\delta = \sigma\epsilon$ for some substitution ϵ . Since $L_1\delta \in \text{gnd}(L_1; \pi_1)$, $\pi_1\sigma\epsilon$ must be true ($i = 1, 2$). But then $L_1\delta \geq L_1\sigma$, and $L_1\delta = (L_1\sigma)\epsilon \in \text{gnd}(L_1\sigma; \pi_1\sigma \wedge \pi_2\sigma)$ both hold.

(\supseteq): Now, assume that $(L_1\sigma)\epsilon$ is a literal from $\text{gnd}(L_1\sigma; \pi_1\sigma \wedge \pi_2\sigma)$. Then, since σ is the most general unifier, we know that $L_1\sigma\epsilon \geq L_i$ hold for both $i = 1, 2$. Furthermore, $\pi_i\sigma\epsilon$ is true ($i = 1, 2$), and thus, $L_1\sigma\epsilon \in \text{gnd}(L_1; \pi_1)$ and $L_1\sigma\epsilon = L_2\sigma\epsilon \in \text{gnd}(L_2; \pi_2)$. **Qed.**

We note that the case when no unifier exists is trivial.

Example 2.13 Consider the following constrained literals

- $(L; \pi) = (P(x, y); (x, y) \neq (v, v) \wedge x \neq a \wedge y \neq b)$
- $(L'; \pi') = (P(z, a); z \neq b)$

Then according to the definition above

$$(L; \pi) \wedge (L'; \pi') = (P(z, a); (z, a) \neq (v, v) \wedge z \neq a \wedge a \neq b \wedge z \neq b)$$

which can be simplified to

$$(P(z, a); z \neq a \wedge z \neq b)$$

This expression is empty over $\mathcal{D}_2 = \{a, b\}$, and covers exactly the atom $P(c, a)$ over $\mathcal{D}_3 = \{a, b, c\}$.

Difference

The *difference*, or *relative difference*, $(L; \pi)$ of two constrained literals $(L_1; \pi_1)$, $(L_2; \pi_2)$ satisfies

$$\text{gnd}(L; \pi) = \text{gnd}(L_1; \pi_1) - \text{gnd}(L_2; \pi_2)$$

Again, if the two literals are unifiable, such a π does exist for any finite domain - in the worst case we just add ground constraints to rule out the disallowed atoms. However, this operation might increase the size of π exponentially, as demonstrated by the example below.

Example 2.14 Consider the difference

$$(L(x_1, x_2, x_3); \top) - (L(x_1, x_2, x_3); \bigwedge_{i=1}^3 x_i \neq a)$$

where $\text{arity}(L) = 3$. If $\mathcal{D}_2 = \{a, b\}$, we might get the still simple expression

$$(L(x_1, x_2, x_3); (x_1, x_2, x_3) \neq (b, b, b))$$

However, if $\mathcal{D}_3 = \{a, b, c\}$, the best we can get is

$$(L(x_1, x_2, x_3); (x_1, x_2, x_3) \neq (b, b, b) \wedge (x_1, x_2, x_3) \neq (c, b, b) \wedge \dots \neq (c, c, c))$$

It is easy to see that in general, if $|\mathcal{D}| = n$ with $a \in \mathcal{D}$, and $\text{arity}(L) = r$, the size of the resulting constraint is $O((n-1)^r)$.

Alternatively, one might take a set of disjoint constrained literals describing the difference as follows. First, take the simpler case when L_1 and L_2 are the same literal L , and consider the difference $(L; \pi_1) - (L; \pi_2)$. Assume $\pi_1 = \bigwedge_{i \in \mathcal{I}_1} \nu_i$, $\pi_2 = \bigwedge_{i \in \mathcal{I}_2} \eta_i$, and $\{\sigma_i \mid i \in \mathcal{I}_2\}$ is the set of induced substitutions for π_2 . Then, the constrained literal set

$$\{(L\sigma_i; \pi_1\sigma_i) \mid i \in \mathcal{I}_2\}$$

describes the difference, i.e.

Lemma 2.15

$$\bigcup_{i \in \mathcal{I}_2} \text{gnd}(L\sigma_i; \pi_1\sigma_i) = \text{gnd}(L; \pi_1) - \text{gnd}(L; \pi_2)$$

Proof:

(\supseteq): Assume $L\delta \in (\text{gnd}(L; \pi_1) - \text{gnd}(L; \pi_2))$. Since $L\delta \notin \text{gnd}(L; \pi_2)$, a subconstraint $\eta_i \in \pi_2$ must be violated, i.e. for some $i \in \mathcal{I}_2$, $\eta_i\delta = \perp$.

Then, by the earlier Remark 2.6, $\delta \geq \sigma_i$ where σ_i is the corresponding induced substitution. Thus, $\delta = \sigma_i\xi$ for some substitution ξ . Finally, since $\pi_1\delta = (\pi_1\sigma_i)\xi$, $L\delta \in \text{gnd}(L\sigma_i; \pi_1\sigma_i)$ must hold.

(\subseteq): Now, assume $(L\sigma_i)\xi \in \text{gnd}(L\sigma_i; \pi_1\sigma_i)$ for some $i \in \mathcal{I}_2$ and grounding substitution ξ . Then, we know that $\pi_1\sigma_i\xi = \top$, and that $\pi_2\sigma_i\xi = \perp$ since $\eta_i\sigma_i\xi = \perp$. Thus, $L\sigma_i\xi \in (\text{gnd}(L; \pi_1) - \text{gnd}(L; \pi_2))$. **Qed.**

However, this set is not pairwise disjoint, and therefore a further step is needed for our purposes.

Lemma 2.16 *W.l.o.g. assume $\mathcal{I}_2 = \{1, \dots, l\}$, and take*

$$\{(L\sigma_i; \pi_1\sigma_i \wedge \eta_1\sigma_i \wedge \eta_2\sigma_i \wedge \dots \wedge \eta_{i-1}\sigma_i) \mid i = 1, \dots, l\}$$

Then this set still describes the difference and its elements are pairwise disjoint.

Proof: We only prove one inclusion, as the other direction is analogous to the first proof, and disjointness trivially follows from the definition of the set.

(\supseteq): Assume $L\delta$ is a ground literal from the difference. Thus, $\pi_1\delta = \top$ and $\eta_i\delta = \perp$ for at least one $i \in \mathcal{I}_2$. Let i be the smallest (left-most) such index.

Then, $\delta \geq \sigma_i$ must hold along with $\eta_j\delta = \top$ for each $j < i$ from \mathcal{I}_2 . Thus, $L\delta \in \text{gnd}(L\sigma_i; \pi_1\sigma_i \wedge \eta_1\sigma_i \wedge \eta_2\sigma_i \wedge \dots \wedge \eta_{i-1}\sigma_i)$. **Qed.**

We also note that the above manipulations preserve the variable disjointness of the left-hand and right-hand sides.

Example 2.17 *Carrying on with example 2.14 above, we have $\mathcal{I}_2 = \{1, 2, 3\}$, $\eta_i : x_i \neq a$ and $\sigma_i = \{x_i \leftarrow a\}$, which gives*

$$\{(L(a, x_2, x_3); \top), (L(x_1, a, x_3); x_1 \neq a), (L(x_1, x_2, a); x_1 \neq a \wedge x_2 \neq a)\}$$

as a result.

Let $|\pi|$ denote the size of π . Then, this operation introduces $O(|\mathcal{I}_2|)$ atoms with a maximal constraint size of $O(|\pi_1| + |\pi_2|)$ in general. This gives a total size of $O(2|\pi^*|^2)$ where $|\pi^*| = \max\{|\pi_1|, |\pi_2|\}$. Clearly, it is independent of the domain size.

Lemma 2.18 *Finally, if $L_1 \neq L_2$, but $\text{mgu}(L_1, L_2) = \delta$ exists (otherwise the difference is $(L_1; \pi_1)$), and \vec{u} denotes the argument of the top symbol in L_1 , we get the desired set by adding*

$$(L_1; \pi_1 \wedge \vec{u} \neq \vec{u}\delta\rho)$$

to the set $(L_1\delta; \pi_1\delta) - (L_2\delta; \pi_2\delta)$ where the variable renaming ρ introduces fresh variables for the variables in $\vec{u}\delta$.

Proof: The literals in $\text{gnd}(L_1; \pi_1)$ can be divided into two disjoint group based on whether they are instances of L_2 or not.

Those that are no instances of L_2 are covered by the proposed constrained literal $(L_1; \pi_1 \wedge \vec{u} \neq \vec{u}\delta\rho)$. Clearly, each such literal is in the difference.

The common instances are covered by $(L_1\delta; \pi_1\delta)$. From these literals we have to remove those which are covered by $(L_2; \pi_2)$ as well. Clearly, it is enough to compute the difference $(L_1\delta; \pi_1\delta) - (L_2\delta; \pi_2\delta)$.

The resulting set together with $(L_1; \pi_1 \wedge \vec{u} \neq \vec{u}\delta\rho)$ covers exactly the elements of the difference. **Qed.**

It is easy to see that the proofs above hold even if some of the constraints are the constants \perp or \top , and our definition of induced substitutions in the case of constants supports the proofs.

The disadvantage of the second method is the fragmentation of the constrained literal, especially since after every time we derive a new assignment we have to subtract each unifiable already defined literal.

Remark 2.19 *Whenever we compute a difference and get a set of literals as a result, we carry on working with the literals separately.*

We could extend our constraint language to handle a set of constrained literals as a single expression. In the literature the corresponding constraints are called disjunctive implicit generalizations, see e.g. [33] for details.

Checking Emptiness

The problem of deciding whether a constrained literal is empty. As we mentioned before, this is equivalent with the unsatisfiability of the corresponding constraint.

This is in general a co-NP-complete problem [11][24]. Lassez and Marriott proposes an algorithm for computing explicit representation in [24], which can be used for determining emptiness as well. Their algorithm is based on generating disjoint partitions of instances by instantiating a single variable with every possible function symbol at every step.

We note that the operation is indeed complex, but so is checking subsumption and subsumption resolution in first-order theorem provers, and even *iProver* calls CDCL iteratively. Yet, these techniques are efficient in practice, which we consider an indication that an efficient implementation of NRCL is possible.

We propose here an enumeration-based algorithm. Assume $\mathcal{D} = \{a_1, \dots, a_n\}$ is ordered by $>$, and $a_n > a_{n-1} > \dots > a_0$. For a constrained literal $(L; \pi)$ with left-hand side variables x_1, \dots, x_k , we find a solution (c_1, c_2, \dots, c_k) denoting $\{x_j \leftarrow c_j \mid j = 1, \dots, k\}$ by enumerating the possible assignments starting with (a_0, a_0, \dots, a_0) .

If for an intermediate assignment (c_1, c_2, \dots, c_k) the subconstraint $\vec{x} \neq \vec{t} \in \pi$ is false, then we increase the *value* of the right-most position involved in \vec{x} . If it is already a_n , we reset it to a_0 and increase the next involved variable to the left.

If no further increase is possible, there is no solution. If we get a solution for $\vec{x} \neq \vec{t}$, we pick the left-most involved variable which we changed, and reset all non- \vec{x} variables to a_0 .

By repeating the above steps, we either get a solution satisfying π , or attempt to increase beyond (a_n, a_n, \dots, a_n) , proving the unsatisfiability of the constraint.

Besides simplicity, this algorithm also has the advantage that the solution might be reusable in operations. We only need to make sure that the solutions for the operands are comparable in the sense that they denote the minimal solutions of the respective constraints w.r.t. the same ordering over the possible assignments.

This can be ensured by ordering the variables based on their left-most occurrence. This way, the solutions of any two non-empty constrained literals with the same predicate symbol are comparable by taking the arguments of the ground literals representing the solutions.

Then, it is enough to consider substitution and adding new subconstraints as primitive operations. Both meet and difference builds upon these steps.

When applying a substitution σ , we check if the current solution satisfies the positive equality constraints induced by σ . If yes, then we apply the substitution and keep this solution.

Otherwise, we keep enumerating by always checking the positive conditions first. If we find an assignment satisfying both π and σ , we apply σ and save the new solutions. If no solutions to be found, the new constrained literal is empty.

When extending π with a new subconstraint $\vec{x} \neq \vec{t}$, we simply continue the enumeration with the current solution and the extended constraint $\pi \wedge \vec{x} \neq \vec{t}$.

2.4 Model Representation

Model Candidate

On the course of this paper, we represent a *model candidate*, also called a *model assumption*, as a set Γ of constrained literals.

Definition 2.20 *A set of constrained literals Γ is called consistent if there is no ground atom covered by both a positive and a negative literal from Γ .*

Γ is strongly consistent if its elements are pairwise disjoint w.r.t. covering atoms, i.e. for all different $(L; \pi), (L'; \pi') \in \Gamma$, $\text{gnd}(|L|; \pi) \cap \text{gnd}(|L'|; \pi') = \emptyset$.

We consider only strongly consistent sets in this paper.

Definition 2.21 (Induced Interpretation) *The set of positive constrained literals in Γ is denoted by Γ^+ . Then the first-order interpretation I_Γ induced by Γ is given as*

$$I_\Gamma = \bigcup_{(L; \pi) \in \Gamma^+} \text{gnd}(L; \pi)$$

This interpretation serves as a minimal model defined by the positive literals, and it is used in the rule *Success* and the relevant proofs.

Trail

NRCL attempts to lift the classic CDCL, and as such, it uses a sequence of literals to store the current partial model assumption.

This *trail* in our case is a sequence of annotated constrained literals. We retain the notation Γ , and extend all our definitions and operations for sets of constrained literals to trails as well. We call the elements of Γ *assignments*, as they define truth-values of ground atoms.

Literals in Γ are either *decision* or *deduced literals*. Decisions are annotated with a unique positive integer, with $(L; \pi)^i$ representing the i th decision in Γ . Deduced literals are annotated with their reasons, a first-order clause from the current clause set. In the course of the paper, α is used to denote an arbitrary annotation, C to denote a reason clause, and k, l, i to denote integers.

We define the value of a ground literal or ground clause *true*, *false*, or *undefined* under Γ lifting the notions of CDCL. In particular, a ground literal L' is *defined by* a constrained literal $(L; \pi) \in \Gamma$, iff $|L'| \in \text{gnd}(|L|; \pi)$. If such an $(L; \pi)$ exists, we also say that Γ *defines* L' . Then, the value of the defined ground literal L' is *true* iff L' and L has the same polarity.

Non-ground literals are treated as unit clauses, and a set of ground clauses represented by the constrained clause $(C; \pi)$ is *true* or *false* in Γ , if all of the covered ground instances are *true*, or *false*, respectively. The notion of *defined by* Γ extends to constrained clauses similarly, i.e. $(C; \sigma)$ is defined w.r.t. Γ iff for each $C' \in \text{gnd}(C; \sigma)$, at least one $L' \in C'$ is defined under Γ . We note that the definition of a *false* constrained clause is non-standard, and it is formulated this way to conveniently define later the invariants and the rule *Conflict*. The notion of *defined by* Γ can be extended to constrained clauses similarly.

The level $\text{lvl}(L)$ of a ground literal L w.r.t. the trail Γ is defined as in CDCL: The annotation of a decision in Γ is the *level* of this decision literal. Then, the level $\text{lvl}(L)$ of a defined ground literal L w.r.t. Γ is the level of the last decision in Γ before the constrained literal defining L , and zero if no such decision exists.

If k is the level of a literal, we might also say *the literal is of level k* . We call the largest level occurring in a trail the *top-level*, and also the level of the trail. If no decision occurs in the trail, it is considered 0.

Following the terminology of SAT solvers, we call a ground clause *assertive* iff it is false w.r.t. the current trail and contains exactly one top-level literal.

Finally, we say a first-order clause C or a constrained clause $(C; \pi)$ is *assertive* iff $\text{gnd}(C)$, and $\text{gnd}(C; \pi)$ contains at least one assertive ground clause, respectively.

Induced Abstraction

Using Γ to define truth-values for groups of ground atoms represented by constrained literals can also be seen as providing a propositional abstraction and an abstract partial interpretation.

In this context, our calculus can be seen as a fine-grained abstraction-refinement algorithm, which interleaves refinement and abstract model search, and lets the clauses and decision heuristics guide the implicit abstraction and refinement steps.

Below, we provide the related definitions and use these later to define our induced ordering. Beyond this, we do not take any advantage of this connection. Further investigation this direction and utilizing existing results for abstraction-refinement-based procedures is left for future work.

We call a set of positive constrained literals Φ an *abstraction*. An abstraction Φ provides a (partial) partitioning of \mathcal{A}_Σ , and by identifying its elements with propositional atoms, we can assign a propositional abstraction to our clause set N .

These propositional atoms are called *abstract atoms*. The notions *abstract literal* and *abstract clause* are the corresponding syntactic expressions built from

abstract atoms. We use the *abstraction function* def_Φ , or simply def , to assign the set of abstract expressions to literals or clauses w.r.t. an abstraction Φ .

Then an *abstract interpretation* over an abstraction Φ is a propositional interpretation over the corresponding abstract atoms.

If the totality of def_Φ is needed, we identify uncovered ground atoms with the unique abstract atom \perp , and the domain of the interpretation is extended accordingly.

The abstraction Φ_Γ induced by Γ is defined as

$$\Phi_\Gamma = |\Gamma| = \{(|L|; \pi) \mid (L; \pi) \in \Gamma\}$$

If Γ is strongly consistent, Φ_Γ is always consistent, i.e. any ground atom is covered by at most one element of Φ .

Γ can be seen as defining an abstract interpretation over Φ_Γ assigning truth-values to abstract atoms based on the polarity of the corresponding constrained literals in Γ , and *undef* to the abstract atom \perp .

2.5 Induced Ordering

In the following, let $<$ denote a given well-founded total ordering over ground expressions - atoms, literals and clauses. Furthermore, let Γ denote a strongly consistent trail.

Definition 2.22 The abstraction function def defined by Γ is given as

$$\text{def}(P) = \begin{cases} (L; \pi) & \text{if } (L; \pi) \in \Gamma \text{ and } (L; \pi) \text{ defines } P \\ \perp & \text{if no such } (L; \pi) \in \Gamma \text{ exists} \end{cases}$$

for each $P \in \mathcal{A}_\Sigma$.

Then, def can be extended to ground literals and clauses by assigning the corresponding negated abstract atom to a negative literal, and the disjunction of the corresponding abstract literals to a clause, respectively.

Definition 2.23 The precedence ordering $<_p^\Gamma$ ($<_p$) defined by Γ is the ordering over the constrained literals in Γ defined by their position in Γ , i.e.

$$(L_1; \pi_1) <_p (L_2; \pi_2) \text{ iff}$$

$$\Gamma = \Gamma_1, (L_1; \pi_1)^{\alpha_1}, \Gamma_2, (L_2; \pi_2)^{\alpha_2}, \Gamma_3$$

for some $\Gamma_1, \Gamma_2, \Gamma_3$ and annotations α_1, α_2 .

We extend the ordering to $\Gamma \cup \{\perp\}$ with \perp as maximal element. Finally, this ordering is extended to abstract literals and clauses as usual.

Definition 2.24 The ordering $<_\Gamma^{\text{atom}}$ induced by Γ is defined over \mathcal{A}_Σ and given as follows: $P <_\Gamma^{\text{atom}} Q$ iff either

1. $\text{def}(P) <_p \text{def}(Q)$, or
2. $\text{def}(P) = \text{def}(Q)$ and $P < Q$

The ordering is extended to ground literals in the usual way, resulting in the literal ordering $<_\Gamma^{\text{lit}}$.

Finally, we extend it to ground clauses: $C <_\Gamma C'$ iff either

1. $\text{def}(C) <_p \text{def}(C')$, or
2. $\text{def}(C) = \text{def}(C')$ and $C (<_{\Gamma}^{lit})_{mul} C'$

where $(<_{\Gamma}^{lit})_{mul}$ denotes the multiset extension of the literal ordering.

$<_{\Gamma}$ extends the atom and literal orderings, and we call it the ordering induced by Γ .

Proposition 2.25 $<_{\Gamma}$ is well-defined, total on ground clauses, and a well-founded ordering.

Proof: It is easy to see that both $<_p$ and $(<_{\Gamma}^{lit})_{mul}$ are well-founded and total orderings over ground clauses. Since $<_{\Gamma}$ is the lexicographical combination of these orderings, $<_{\Gamma}$ inherits these properties. **Qed.**

This dynamic ordering captures the local correlation between the atoms and literals in the search, and shifts the focus on the recent behavior of the calculus. Finally, we introduce an easily provable proposition, which is used in the proofs later on.

Proposition 2.26 Let $<$ be an arbitrary well-founded and total ordering over ground clauses, S and S' finite sets of ground clauses, and assume there is a function $\gamma : S' \rightarrow S$ such that for each $C \in S'$, $C < \gamma(C)$.

Then $S' < S$ holds w.r.t. the multiset extension of $<$.

3 Calculus

The calculus NRCL attempts to find a model through a series of both arbitrary and deduced assignments. Analogous to the propositional SAT solvers, we apply propagation to find literals implied by existing assignments, and once it is exhausted, we add arbitrary literals, so-called *decisions* to the trail.

We call this phase *conflict search* and it ends with either a model of the original clause set, or with finding a clause C with some instances given in the form $(C; \sigma; \pi)$ falsified by the current trail. In the latter case, we start *conflict resolution* and through resolving the current false clause with reason clauses from the trail, we learn a new assertive clause and backtrack to a state where this clause is not yet falsified by removing some of the most recent assignments from the trail.

As opposed to propositional SAT solving, where every clause can be considered already exhaustively factorized, in our case some ground instances might be still subject to factorization, and this requires further rules.

The rule *Factorize* handles this during clause learning. However, the calculus might still reach a state where the right-most literal on the trail is the last decision, the learnable clause is not assertive, but no factorization is possible. When such a state is reached, we simply learn the current candidate for clause learning. To avoid this situation again, we further demand that a new decision should not falsify any clause instance immediately, unless *Factorize* is applicable.

We call a clause *blocking* a new decision if adding the decision to the trail would falsify an instance of the clause without allowing *Factorize* to handle the immediate conflict, see the precise definition below. We note that a clause learned in the above fashion blocks the last decision.

Definition 3.1 We say that a decision $(L; \pi)$ is blocked in Γ by a clause C , if C has a ground instance $C\sigma$ with $L_1, L_2 \in C\sigma$ such that for $\Gamma' = \Gamma, (L; \pi)$

- $C\sigma$ is false under Γ'
- $(L; \pi)$ is undefined in Γ
- L_1 and L_2 become false by the decision, i.e. $\neg L_1, \neg L_2 \in \text{gnd}(L; \pi)$
- $L_1 \neq L_2$

If no such C exists, we say that the decision is not blocked in Γ .

Example 3.2 Consider $\mathcal{D} = \{a, b, c\}$, $\Gamma = \{(\neg Q(x, y); \top)^1\}$, and

$$N = \{C : \neg P(x) \vee \neg P(y) \vee Q(x, y), \dots\}$$

Then the decisions $(P(x); \top)$, $(P(x); x \neq c)$ are both blocked in Γ by C , as witnessed by the ground instance $\neg P(a) \vee \neg P(b) \vee Q(a, b)$.

We give our calculus as a set of rules over so-called *states*, tuples of the form

$$(\Gamma; N; U; k; s)$$

where Γ denotes the trail, N the given clause set, U the set of learned clauses, k a non-negative integer - unless terminating with *Success* -, and s a *state indicator*. The latter can be \top , \perp , or a set of clause instances $\text{gnd}(C\sigma; \pi)$ given as $(C; \sigma; \pi)$.

\top indicates the conflict search phase, if $k \geq 0$, or that Γ defines a model for N , if $k = -1$. \perp means the empty clause has been learned, i.e. the unsatisfiability of N has been established. Finally, an indicator of the form $(C; \sigma; \pi)$ represents a set of clause instances falsified by the current trail Γ , and indicates the conflict resolution phase of our calculus.

Our results extend to any derivation starting from a *sound* state (see Definition 4.1). Here we propose the initial state

$$(\epsilon; N; \emptyset; 0; \top)$$

where ϵ stands for the empty trail, and N is the set of input clauses.

Next, we address a technical question regarding deduced literals and conflict resolution. It is often the case that for a clause $(C \vee L) \in N$ with $C\sigma$ implying $(L\sigma; \pi)$ for some σ, π w.r.t. the current trail Γ , the involved substitution σ substitutes variables not occurring in L . See Example 3.3 for a demonstration of this behavior.

Should we save only $(L\sigma; \pi)$ to the trail, we would lose this part of the assignment. However, during conflict resolution we need the exact clause instances responsible for the assignment. Therefore, to avoid recomputing the relevant substitutions, we save the *constrained closure* $(L \cdot \sigma; \pi)$, where $L \cdot \sigma$ is the *closure* representing $L\sigma$.

This is an extension of the existing notation for the sake of clause learning. For all other purposes, $L \cdot \sigma$ is identified with $L\sigma$, and all definitions over constrained literals can be extended to constrained closures accordingly. The literal L is also considered to be a short-hand for $L \cdot \emptyset$. We also note that in our calculus decisions are always considered having empty closures.

Finally, a short remark on the usage of the operations over constrained expressions: Conjunction is used whenever we try to unify two constrained literals, e.g. during learning a new clause via resolution, or finding candidates for propagation. Difference is needed when we remove already defined literals ensuring that a new assignment only defines new values. Emptiness is tested overall in the calculus to ensure that a new assignment indeed defines the value of at least one ground atom.

Below, we provide the rules of our calculus in a generic style as a state transition system, similarly to [27]. We note that in the rules π_1, π_2 is often used as a short-hand for $\pi_1 \wedge \pi_2$, if it is unambiguous. Furthermore, blocking is considered only w.r.t. the current clause set $N \cup U$ in the rest of the paper. For further details on the applied strategy and technicalities, see Section 5 and Section 8.

3.1 Rules for Conflict Search

Propagate

$$(\Gamma; N; U; k; \top) \Rightarrow (\Gamma, (L \cdot \sigma; \pi)^{C \vee L}; N; U; k; \top)$$

if $k \geq 0$, and for $(C \vee L) \in (N \cup U)$, σ , and π

- $(C\sigma; \pi)$ is false under Γ
- $(L\sigma; \pi)$ is undefined in Γ
- $(L\sigma; \pi)$ is not empty

This rule deduces new literals which have to be true under the current model assumption. The conditions ensure that this step is sound and *effective*, i.e. each ground literal defined by the added literal is indeed a consequence and at least one such literal exists.

Example 3.3 Let $a, b \in \mathcal{D}$, N and U arbitrary, $C_1, C_2 \in N$, and the current state

$$(\Gamma; N; U; 1; \top)$$

where

$$\Gamma = (P(x, x); \top)^{C_1}, (Q(a, x); \top)^{C_2}, (\neg P(x, y); (x, y) \neq (v, v))^1$$

Then, if $C = P(y, b) \vee \neg Q(x, y) \vee R(y)$ is a clause from N , Propagate can be applied for C , and we might get the state

$$(\Gamma, (R(y) \cdot \{x \leftarrow a\}; y \neq b)^C; N; U; 1; \top)$$

Decide

$$(\Gamma; N; U; k; \top) \Rightarrow (\Gamma, (L; \pi)^{k+1}; N; U; k+1; \top)$$

if $k \geq 0$, and for L, π

- $(L; \pi)$ is undefined in Γ
- $(L; \pi)$ is not blocked in Γ
- $(L; \pi)$ is not empty

- $\exists(C \vee L') \in N$ such that $|L| \geq |L'|$, i.e. $\exists \delta : L = L'\delta$, or $L = \neg L'\delta$

Decide adds an assumption to Γ which is not blocked by any of the clauses, and which is effective.

We note that the last condition is optional, it does not influence any of our results. This restriction allows earlier termination with *Success* and keeps the calculus from defining irrelevant ground atoms. After terminating with *Success*, every undefined ground atom can be considered having arbitrary truth-values, or simply *false*, the way it is defined in I_Γ .

We also note that blocking only identifies one kind of immediate conflicts, we might still get to an outright conflict if it can be handled with factorization, see Example 3.4 below, and Lemma 5.5 for details.

Example 3.4 Let $\mathcal{D} = \{a, b, c\}$, $\Gamma = (P(x, x); \top)^{P(x, x)}, (Q(x, a); \top)^{Q(x, a)}$, and

$$N = \{P(x, x), Q(x, a), \neg Q(x, y) \vee P(x, y) \vee P(x, y)\}$$

Then, the decision $(\neg P(x, y); (x, y) \neq (v, v))$ is not blocked, yet

$$(\neg Q(x, y) \vee P(x, y) \vee P(x, y); \{y \leftarrow a\}; x \neq a)$$

is false w.r.t. $\Gamma, (\neg P(x, y); (x, y) \neq (v, v))$ ¹. We note that conflict resolution learns the clause $\neg Q(x, y) \vee P(x, y)$ from this conflict.

We also note that whenever a decision is blocked, we can always pick a stricter unblocked decision, shown below.

Proposition 3.5 For every blocked decision $(L; \pi)$ and blocking clause C , there is a decision $(L\sigma; \pi\sigma, \pi')$ for some σ, π' such that it is not blocked by C and it is not empty.

Proof: It is easy to see that any ground literal from $\text{gnd}(L; \pi)$ satisfies this condition. **Qed.**

Conflict

$$(\Gamma; N; U; k; \top) \Rightarrow (\Gamma; N; U; k; (C; \sigma; \pi))$$

if $k \geq 0$, and for some $\perp \neq C \in (N \cup U)$, σ , and π

- $(C\sigma; \pi)$ is false under Γ
- $(C\sigma; \pi)$ is not empty

Conflict identifies a set of clause instances contradicting the current model assumption. We also refer to this set as the *conflict-set*.

Example 3.6 Let $\mathcal{D} = \{a, b, c\}$, and

$$\begin{aligned} N = \{ & C_1 : \neg P(c), C_2 : \neg P(x) \vee \neg P(y) \vee Q(x, y), \\ & C_3 : \neg P(y) \vee \neg Q(a, y), C_4 : \neg Q(x, b) \vee \neg P(x) \} \\ \Gamma = & (\neg P(c); \top)^{C_1}, (P(x); x \neq c)^1, (\neg Q(a, y); y \neq c)^{C_3} \end{aligned}$$

Then the following is a valid step:

$$(\Gamma; N; \emptyset; 1; \top) \xrightarrow{\text{Conflict}(C_2)} (\Gamma; N; \emptyset; 1; (\neg P(x) \vee \neg P(y) \vee Q(x, y); \{x \leftarrow a\}; y \neq c))$$

Success

$$(\Gamma; N; U; k; \top) \Rightarrow (\Gamma; N; U; -1; \top)$$

if $k \geq 0$, and $I_\Gamma \models N$.

We note that the last condition, $I_\Gamma \models N$, can be replaced by demanding that the rules *Propagate*, *Decide* and *Conflict* are exhausted and $\perp \notin (N \cup U)$.

From this it follows that each ground atom is defined and there is no falsified instance, i.e. every ground clause $C \in \text{gnd}(N \cup U)$ is true w.r.t. the current trail.

Failure

$$(\Gamma; N; U; k; \top) \Rightarrow (\Gamma; N; U; 0; \perp)$$

if $\perp \in (N \cup U)$.

The two terminal rules correspond to the satisfiability and unsatisfiability of the clause set, respectively. Unsatisfiability is detected through learning the empty clause \perp .

3.2 Rules for Conflict Resolution**Skip**

$$(\Gamma, (L' \cdot \sigma'; \pi')^{C'}; N; U; k; (C; \sigma; \pi)) \Rightarrow (\Gamma; N; U; k; (C; \sigma; \pi))$$

if there is no $L \in C$ such that

- $\exists \eta = \text{mgu}(L' \sigma', \neg L \sigma)$, and
- $(C \sigma \eta; \pi \eta, \pi' \eta)$ is not empty

Skip drops the right-most literal from the trail during conflict resolution if it is not a decision and it does not contribute to the conflict, i.e. it does not touch any instance of the conflict-set.

Resolve

$$(\Gamma, (L' \cdot \sigma'; \pi')^{C' \vee L'}; N; U; k; (C \vee L; \sigma; \pi)) \Rightarrow (\Gamma, (L' \cdot \sigma'; \pi')^{C' \vee L'}; N; U; k; ((C \vee C')_{\eta_0}; \sigma^*; \pi \eta, \pi' \eta))$$

if for L' , σ , π' and $C' \vee L'$, and

- $((C \vee L) \sigma; \pi)$ is not assertive, or $k = 0$
- $\exists \eta = \text{mgu}(L' \sigma', \neg L \sigma)$, and let
 - $\eta_0 = \text{mgu}(L', \neg L)$
 - σ^* such that $\sigma \sigma' \eta = \eta_0 \sigma^*$
- $((C \vee L) \sigma \eta; \pi \eta, \pi' \eta)$ is not empty

We note that keeping $\sigma^*|_{\text{var}((C \vee C')_{\eta_0})}$ instead of σ^* is enough for the soundness of the rule and our calculus, as it contains all the relevant information. Furthermore, the existence of η implies the existence of η_0 and σ^* .

If the right-most literal in Γ is not a decision and is involved in the conflict-set, we proceed with resolution. The conditions imply that there are corresponding ground inferences and the new conflict-set is not empty.

Note that dropping the used literal is not desired as the new conflict might still be resolvable with it.

Factorize

$$(\Gamma, \ell; N; U; k; (C \vee L_1 \vee L_2; \sigma; \pi)) \Rightarrow (\Gamma, \ell; N; U; k; ((C \vee L_1)\eta_0; \sigma^*; \pi\eta))$$

if $\ell = (L' \cdot \sigma'; \pi')^\alpha$ for some L' , σ' , π' , and annotation α , and

- $\exists \eta = \text{mgu}\{L_1\sigma, L_2\sigma, L'\sigma'\}$, and let
 - $\eta_0 = \text{mgu}(L_1, L_2)$
 - σ^* such that $\sigma\eta = \eta_0\sigma^*$
- $((C \vee L_1)\sigma\eta; \pi\eta, \pi'\eta)$ is not empty

Again, the existence of η implies the existence of η_0 and the appropriate σ^* , and keeping $\sigma^* \upharpoonright_{\text{var}((C \vee L_1)\eta_0)}$ is sufficient. We also note that α can be both a reason clause and a decision level.

Factorize factorizes some of the conflicting ground clauses. As in the case of *Resolve*, the used literal should not be dropped from the trail.

Backjump

$$(\Gamma_1, \Gamma_2; N; U; k; (C; \sigma; \pi)) \Rightarrow (\Gamma_1; N; U \cup \{C\}; k'; \top)$$

if $0 \leq k' \leq k$, $k' = \text{lvl}(\Gamma_1)$, and one of the following condition-sets hold:

- (1) $k = 0$, and $C = \perp$, or
- (2) $k > 0$, $(C\sigma; \pi)$ is assertive, and C has no false instance under Γ_1 , or
- (3) $k > 0$, the right-most element of Γ_2 is the top-level decision, $(C\sigma; \pi)$ is not assertive, *Factorize* cannot be applied, and C has no false instance under Γ_1

It is clear that $k' = 0$ or $k' < k$ in case (1) and (2), (3), respectively.

The optimal choice for k' is the smallest level for which the learned clause can be used in *Propagate*. Such a k' might not always exist for the learned clause C , largely due to the instances of C not covered by $(C\sigma; \pi)$. In these cases the optimal choice for k' is the largest level for which C has no false instance. For more details see Section 8.

In case (1), we say that *the empty clause \perp is learned*. In case (2), we say *a new assertive clause is learned*, and in case (3) *a new blocking clause is learned*.

The latter clause is indeed blocking the last decision under some regularity conditions, see Lemma 5.4 for details. We note that case (3) can indeed occur as the following example demonstrates:

Example 3.7 (Learning a blocking clause)

Consider the clause set

$$N = \{C_1 : R(x, x), C_2 : P(x) \vee \neg Q(x, y), C_3 : R(x, y) \vee Q(x, y) \vee P(x) \vee P(y)\}$$

and let $\Gamma = \Gamma', (\neg Q(x, y); \top)^{C_2}$ with

$$\Gamma' = (R(x, x); \top)^{C_1}, (\neg R(x, y); (x, y) \neq (v, v))^1, (\neg P(x); \top)^2$$

Then the following is a valid conflict resolution:

$$\begin{aligned}
& (\Gamma; N; \emptyset; 2; (R(x, y) \vee Q(x, y) \vee P(x) \vee P(y); \emptyset; (x, y) \neq (v, v))) \xrightarrow{\text{Resolve}} \\
& (\Gamma; N; \emptyset; 2; (R(x, y) \vee P(x) \vee P(x) \vee P(y); \emptyset; (x, y) \neq (v, v))) \xrightarrow{\text{Skip}} \\
& (\Gamma'; N; \emptyset; 2; (R(x, y) \vee P(x) \vee P(x) \vee P(y); \emptyset; (x, y) \neq (v, v))) \xrightarrow{\text{Factorize}} \\
& (\Gamma'; N; \emptyset; 2; (R(x, y) \vee P(x) \vee P(y); \emptyset; (x, y) \neq (v, v))) \xrightarrow{\text{Backjump}(3)} \\
& ((R(x, x); \top)^{C_1}, (\neg R(x, y); (x, y) \neq (v, v))^1; N; \{R(x, y) \vee P(x) \vee P(y)\}; 1; \top)
\end{aligned}$$

Remark 3.8 We also wish to note that the current formulation of the calculus handles blocking decisions and learning blocking clauses asymmetrically in the following sense.

Let $\mathcal{D} = \{a, b, c\}$, $N = \{P(x, x), Q(x, a), \neg Q(x, y) \vee P(x, y) \vee P(x', y)\}$, and

$$\Gamma = (P(x, x); \top)^{P(x, x)}, (Q(x, a); \top)^{Q(x, a)}$$

Then the decision $(\neg P(x, y); (x, y) \neq (z, z))$ is blocked by $\neg Q(x, y) \vee P(x, y) \vee P(x', y)$. We could use factorization and learn $(\neg Q(x, y) \vee P(x, y))$, but instead we rather throw away the decision candidate and try another.

On the other hand, if in some regular run (see Definition 5.2) a conflict state of the form

$$(\Gamma', \ell^k; N'; U'; k; (\neg Q(x, y) \vee P(x, y) \vee P(x', y); \{y \leftarrow a\}; x \neq a \wedge x' \neq a))$$

with $\ell = (\neg P(x, y); (x, y) \neq (z, z))$ arises, we choose Factorize over learning a blocking clause outright - there is indeed a blocking instance -, and learn the assertive and not-blocking $\neg Q(x, y) \vee P(x, y)$ in the end.

3.3 Example

Example 3.9 As an example, we present a derivation which constructs a model over $\mathcal{D} = \{a, b, c\}$ for the clause set

$$\begin{aligned}
N = \{ & C_1 : \neg P(c, x, x), C_2 : \neg P(x, y, z) \vee \neg P(u, w, t) \vee Q(x, u), \\
& C_3 : \neg P(x, y, z) \vee \neg Q(a, x), C_4 : \neg Q(x, b) \vee \neg P(x, y, z) \}
\end{aligned}$$

The run below is by no means optimal - any sensible heuristic would choose the negative assignment for P outright -, but it is a valid derivation, and serves well as a demonstration for the syntactic behavior.

$$\begin{aligned}
& (\epsilon; N; \emptyset; 0; \top) \xrightarrow{\text{Propagate}} ((\neg P(c, x, x); \top)^{C_1}; N; \emptyset; 0; \top) \xrightarrow{\text{Decide}} \\
& ((\neg P(c, x, x); \top)^{C_1}, (P(x, y, z); x \neq c)^1; N; \emptyset; 1; \top) \xrightarrow{\text{Propagate}} \\
& ((\neg P(c, x, x); \top)^{C_1}, (P(x, y, z); x \neq c)^1, (\neg Q(a, x); x \neq c)^{C_3}; N; \emptyset; 1; \top) \xrightarrow{\text{Conflict } C_2} \\
& (\dots; N; \emptyset; 1; (\neg P(x, y, z) \vee \neg P(u, w, t) \vee Q(x, u); \{x \leftarrow a\}; u \neq c)) \xrightarrow{\text{Resolve}} \\
& (\dots; N; \emptyset; 1; (\neg P(a, y, z) \vee \neg P(u, w, t) \vee \neg P(u, y', z'); \emptyset; u \neq c)) \xrightarrow{\text{Skip}} \\
& (\dots, (P(x, y, z); x \neq c)^1; N; \emptyset; 1; (\neg P(a, y, z) \vee \neg P(u, w, t) \vee \neg P(u, y', z'); \emptyset; u \neq c)) \\
& \xrightarrow{\text{Factorize}} (\dots, (P(x, y, z); x \neq c)^1; N; \emptyset; 1; (\neg P(a, y, z) \vee \neg P(u, w, t); \emptyset; u \neq c))
\end{aligned}$$

$$\begin{aligned}
& \xRightarrow{\text{Factorize}} ((\neg P(c, x, x); \top)^{C_1}, (P(x, y, z); x \neq c)^1; N; \emptyset; 1; (\neg P(a, y, z); \emptyset; \top)) \\
\text{Let } U_1 = \{C_5 : \neg P(a, y, z)\}. \\
& \xRightarrow{\text{Backjump}(2)} ((\neg P(c, x, x); \top)^{C_1}; N; U_1; 0; \top) \xRightarrow{\text{Propagate}} \\
& ((\neg P(c, x, x); \top)^{C_1}, (\neg P(a, y, z); \top)^{C_5}; N; U_1; 0; \top) \xRightarrow{\text{Decide}} \\
& ((\neg P(c, x, x); \top)^{C_1}, (\neg P(a, y, z); \top)^{C_5}, (P(b, y, z); \top)^1; N; U_1; 1; \top) \xRightarrow{\text{Propagate}} \\
& (\dots, (P(b, y, z); \top)^1, (Q(x, u) \cdot \sigma_1; \top)^{C_2}; N; U_1; 1; \top) \xRightarrow{\text{Conflict}C_4}
\end{aligned}$$

Where $\sigma_1 = \{x \leftarrow b, u \leftarrow b\}$.

$$\begin{aligned}
& (\dots, (Q(x, u) \cdot \sigma_1; \top)^{C_2}; N; U_1; 1; (\neg Q(x, b) \vee \neg P(x, y, z); \{x \leftarrow b\}; \top)) \xRightarrow{\text{Resolve}} \\
& (\dots; N; U_1; 1; (\neg P(x, y, z) \vee \neg P(x, y', z') \vee \neg P(b, w, t); \{x \leftarrow b\}; \top)) \xRightarrow{\text{Skip}} \\
& (\dots, (P(b, y, z); \top)^1; N; U_1; 1; (\neg P(x, y, z) \vee \neg P(x, y', z') \vee \neg P(b, w, t); \{x \leftarrow b\}; \top)) \\
& \xRightarrow{\text{Factorize}} (\dots, (P(b, y, z); \top)^1; N; U_1; 1; (\neg P(x, y, z) \vee \neg P(b, w, t); \{x \leftarrow b\}; \top)) \\
& \xRightarrow{\text{Factorize}} (\dots, (P(b, y, z); \top)^1; N; U_1; 1; (\neg P(b); \emptyset; \top)) \xRightarrow{\text{Backjump}(2)} \\
& ((\neg P(c, x, x); \top)^{C_1}, (\neg P(a, y, z); \top)^{C_5}; N; U_1 \cup \{\neg P(b, y, z)\}; 0; \top) \xRightarrow{\text{Propagate}} \\
& ((\neg P(c, x, x); \top)^{C_1}, (\neg P(a, y, z); \top)^{C_5}, (\neg P(b, y, z); \top)^{C_6}; N; U_2; 0; \top) \xRightarrow{\text{Decide}} \\
\text{Where } U_2 = U_1 \cup \{C_6 : \neg P(b, y, z)\}. \\
& (\dots, (\neg P(b, y, z); \top)^{C_6}, (\neg P(c, y, z); (y, z) \neq (v, v))^1; N; U_2; 1; \top) \xRightarrow{\text{Decide}} \\
& (\dots, (\neg P(c, y, z); (y, z) \neq (v, v))^1, (Q(x, y); \top)^2; N; U_2; 2; \top) \xRightarrow{\text{Success}} \\
& (\dots, (\neg P(c, y, z); (y, z) \neq (v, v))^1, (Q(x, y); \top)^2; N; U_2; -1; \top)
\end{aligned}$$

4 Soundness

Now, we show soundness. The following state invariant defines a consistency notion for states.

Definition 4.1 A state $(\Gamma; N; U; k; s)$ is sound if and only if the followings hold:

1. Γ is a consistent sequence of constrained literals
2. Γ is well-formed, i.e.
 - (a) if $k \geq 0$ then Γ contains exactly k decisions
 - (b) for each i from $1, 2, \dots, k$, there is a unique $(L; \pi)^i \in \Gamma$
 - (c) the decisions occur in Γ in the order of their levels
 - (d) for each decomposition $\Gamma = \Gamma_1, (L; \pi)^i, \Gamma_2; (L, \pi)^i$ satisfies the conditions of Decide w.r.t. Γ_1 , N , and U
 - (e) for each decomposition $\Gamma = \Gamma_1, (L \cdot \sigma; \pi)^{C \vee L}, \Gamma_2; (C\sigma; \pi)$ is false under Γ_1 , and $(L\sigma; \pi)$ satisfies the conditions for Propagate w.r.t. Γ_1 and $C \vee L$
3. $N \models U$
4. $s = \perp$ implies $\perp \in N \cup U$

5. $k = -1$ implies $I_\Gamma \models N$

6. if $s = (C; \sigma; \pi)$ then $(C\sigma; \pi)$ is false under Γ , $N \models C$, and $(C\sigma; \pi)$ is not empty.

A rule is called *sound* iff it preserves the soundness of its left-hand side state.

It is easy to see that the initial state $(\epsilon; N; \emptyset; 0; \top)$ is always sound. Furthermore, soundness is an invariant, since each rule preserves this property, as proven below.

Theorem 4.2 *The rules of NRCL are sound.*

Proof: The soundness of *Propagate*, *Decide*, *Conflict*, and the terminal rules *Failure* and *Success* is straightforward to prove from the definitions themselves, and therefore, we entrust it to the reader.

In the case of *Skip*, dropping the right-most literal $(L' \cdot \sigma'; \pi')^{C'}$ from Γ does preserve the well-formedness and consistency properties of Γ . $N \models U$ remains unchanged and the rest of the conditions are irrelevant in this case, except for the last one.

Now, assume the last property does not hold after applying *Skip*. It is only possible if some ground clause C'' from $\text{gnd}(C\sigma; \pi)$ were false under Γ , $(L' \cdot \sigma'; \pi')^{C'}$, but is undefined under Γ . Thus, $(L' \cdot \sigma'; \pi')$ must have made it false, and therefore, for some δ and $L'' \in C''$, $L'' = \neg L' \cdot \sigma' \delta$ and $\pi' \delta$ is true.

Let L be the literal in C corresponding to L'' . Then, the most general unifier η of $\neg L \sigma$ and $L' \sigma'$ must exist and $C'' \in \text{gnd}(C\sigma\eta; \pi\eta, \pi'\eta)$, which is therefore not empty. This violates the preconditions of *Skip*, a contradiction.

For *Resolve*, it is enough to see that the new clause is a consequence of N , and the new state indicator $((C \vee C')\eta_0; \sigma^*; \pi\eta, \pi'\eta)$ is unsatisfiable under Γ , using the notations of the definition for *Resolve*.

The first claim follows from the soundness of the left-hand side and from the soundness of resolution. As for the second claim, we make the following observations:

- $(C \vee C')\eta_0 \sigma^* = (C' \vee C)\sigma \sigma' \eta$
- Each instance from $\text{gnd}(C' \sigma' \eta; \pi' \eta)$ is false under the current trail, as per the well-formedness conditions for derived literals.
- Each instance from $\text{gnd}(C\sigma\eta; \pi\eta)$ is false under the trail by the soundness of the left-hand side.

From these it follows that each ground clause from $\text{gnd}(C \vee C'; \sigma \sigma' \eta; \pi\eta, \pi' \eta)$ is false under the current trail.

The soundness of *Factorize* can be proven analogously, and the proof for *Backjump* is straightforward. We entrust them to the reader. **Qed.**

Next, we define *runs*, i.e. sound derivations in our calculus.

Definition 4.3 *A run (from a clause set N) is a sequence of states such that each subsequent state is derived with a rule from the previous one, and the initial state is a sound state (with N as the original clause set).*

A direct consequence of Theorem 4.2 is that each state in a run is sound, and in particular, for each conflict resolution state $(\Gamma; N; U; k; (C; \sigma; \pi))$, each ground clause from $\text{gnd}(C\sigma; \pi)$ is false w.r.t. Γ .

Theorem 4.4 (Soundness) *The calculus NRCL is sound, i.e. if a run terminates with the Failure, or Success rules, then the starting set N is unsatisfiable, and satisfiable, respectively. Furthermore, in the latter case the trail upon termination defines a model of N .*

Proof: It follows immediately from the definitions and Theorem 4.2. **Qed.**

5 Regular Runs

In this section, we define a strategy for NRCL in the form of *regular runs*, which is sufficient to prove both non-redundant clause learning, and termination in the later sections.

Definition 5.1 *A sound state $(\Gamma; N; U; k; s)$ is regular iff the following hold:*

- *If $\Gamma = \Gamma', (L \cdot \sigma; \pi)^\alpha$, then no clause from $N \cup U$ is false w.r.t. Γ' .*
- *For all decomposition $\Gamma = \Gamma_1, (L; \pi)^i, \Gamma_2$ with decision $(L; \pi)^i$, Propagate is exhausted w.r.t. Γ_1 and $N \cup U$.*

We note that the last assignment on the trail might still make some clauses false, and the initial state $(\epsilon; N; \emptyset; 0; \top)$ is always regular.

Definition 5.2 *We call a run regular iff the following holds:*

- *The starting state is regular.*
- *During conflict search, rules are always applied in this order exhaustively: terminal rules, Conflict, Propagate, Decide. (Or Failure, Conflict, Propagate, Decide, Success, if we test success through exhausted conflict search.)*
- *In conflict resolution Backjump is always applied as soon as possible, and it backtracks to a regular state.*

Lemma 5.3 *Regular runs preserve regularity, i.e. every state in a regular run is regular.*

Proof: It follows from the definitions, we only note that backjumping to a state which is regular w.r.t. the new learned clause set as well is always possible. If nothing else, the empty trail is always a valid choice. **Qed.**

The backtrack-level proposed in the proof above is not practical, of course. For more details on a more accurate backjumping to a regular state see Section 8.

Below, we show some useful properties of regular runs.

Lemma 5.4 *In a regular run the following hold:*

- (1) *For any deduced literal $(L \cdot \sigma; \pi)^{C \vee L}$ of level k on the trail with $k > 0$, each ground clause in $\text{gnd}((C \vee L)\sigma; \pi)$ contains at least two literals of level k .*

- (2) If $(C; \sigma; \pi)$ represents false clauses in some conflict state, then each ground clause in $\text{gnd}(C\sigma; \pi)$ contains at least two top-level literals, if the state is the result of an application of *Conflict*, and at least one top-level literal otherwise.
- (3) If a clause C is learned according to the case *Backjump*-(3), then it blocks the former top-level decision.

Proof: First, assume $(L \cdot \sigma; \pi)$ is a deduced literal and it was implied by $(C \vee L; \sigma; \pi)$ w.r.t. Γ which was the current trail before the corresponding application of *Propagate*.

Let k be the level of the right-most decision in Γ , and $C' \vee L'$ a ground clause from $\text{gnd}((C \vee L)\sigma; \pi)$ such that L' corresponds to L . Then L' is of level k , of course.

Furthermore, if no other literal in C' is of level k , $C' \vee L'$ would have implied L' before the last decision, which contradicts the exhaustive application of *Propagate*. Thus, $C' \vee L'$ must contain at least two literals of level k .

Second, since conflicts are found immediately, any conflicting non-empty ground clause C' must contain at least one top-level literal. A conflicting ground clause with a single top-level literal, however, would contradict the exhaustive application of *Propagate*. Thus, after applying *Conflict*, all ground clause in the conflict-set contains at least two top-level literals. It only remains to show that the rules *Resolve*, *Skip*, and *Factorize* preserve the weaker property of having at least one top-level literals. Obviously, e.g. *Factorize* can break the stronger property.

We only prove this for *Resolve*, the rest can be shown similarly. Assume that at an application of *Resolve* $(L' \cdot \sigma'; \pi')^{C' \vee L'}$ is the involved deduced literal, $(C \vee \neg L; \sigma; \pi)$ represents the false clauses before, and $((C \vee C')\eta_0; \sigma^*; \pi\eta, \pi'\eta)$ after applying the rule, where $\eta = \text{mgu}(L'\sigma', L\sigma)$, $\eta_0 = \text{mgu}(L', L)$, and σ^* such that $\eta_0\sigma^* = \sigma\sigma'\eta$.

It is easy to see that for every ground clause

$$(C_0 \vee C'_0) \in \text{gnd}((C \vee C')\sigma\sigma'\eta; \pi\eta, \pi'\eta)$$

there are corresponding ground clauses $(C'_0 \vee L'_0) \in \text{gnd}((C' \vee L')\sigma'\eta; \pi'\eta)$ and $(C_0 \vee \neg L_0) \in \text{gnd}((C \vee \neg L)\sigma\eta; \pi\eta)$ whose resolvent is exactly $(C_0 \vee C'_0)$, and L_0, L'_0 correspond to L and L' , respectively, and $L_0 = L'_0$.

Then, by the first claim of this lemma, C'_0 must contain at least one top-level literals, and so does $C_0 \vee C'_0$.

Finally, assume C is learned when case (3) of *Backjump* is applied to the state

$$(\Gamma, (L; \pi)^k; N; U; k; (C; \sigma; \pi))$$

Now, let $(C' \vee L'_1 \vee \dots \vee L'_s) \in \text{gnd}(C\sigma; \pi)$ an arbitrary ground clause, where L'_1, \dots, L'_s denotes the top-level literals of the clause.

By (2), $s \geq 1$, and, since $(C\sigma; \pi)$ has no assertive clause, even $s \geq 2$ must hold. We also know that *Factorize* was not applicable, thus, for any $i \neq j$ from $1, \dots, s$, $L_i \neq L_j$ holds. Thus, C blocks the decision $(L; \pi)$ w.r.t. Γ , as witnessed by the ground clause above. **Qed.**

It can be also shown that if there is an immediate conflict after a decision in a regular run, *Factorize* is applied next.

Lemma 5.5 *Assume*

$$\stackrel{Decide}{\Rightarrow} (\Gamma, (L; \pi)^k; N; U; k; \top) \stackrel{Conflict}{\Rightarrow} (\Gamma, (L; \pi)^k; N; U; k; (C; \sigma'; \pi'))$$

is a valid subderivation in a regular run. Then Factorize, and only Factorize, is applicable to the conflict state $(\Gamma, (L; \pi)^k; N; U; k; (C; \sigma'; \pi'))$.

Proof: Obviously, *Resolve* and *Skip* cannot be applied. Furthermore, if case *Backjump*-(3) were applicable, there would be a ground clause in $\text{gnd}(C\sigma'; \pi')$ blocking the last decision, a contradiction.

Also, there cannot be any ground clause in $\text{gnd}(C\sigma'; \pi')$ with a single top-level literal, since otherwise *Propagate* would not have been applied exhaustively before the decision. And $C = \perp$ cannot hold either, as otherwise *Failure* should have been applied earlier. Thus, the other cases of *Backjump* do not apply either.

Finally, let C_0 a ground clause from $\text{gnd}(C\sigma'; \pi')$. This clause exists, and must contain at least two top-level literals, see Lemma 5.4(2). These literals are falsified by the last decision, and do not block the decision.

Let L_0, K_0 two such literals and $C_0 = C'_0 \vee L_0 \vee K_0$. Then these literals are equal, and the corresponding literals L_1, K_1 in $C\sigma$ are unifiable.

Then *Factorize* is applicable unifying L_1 and K_1 , and $C'_0 \vee L_0$ can be used to prove the non-emptiness condition. **Qed.**

6 Redundancy

We define redundancy w.r.t. the induced ordering $<_\Gamma$ in the standard way:

Definition 6.1 *A ground clause C is redundant w.r.t. a ground clause set N (and $<_\Gamma$) iff*

$$C \in N, \text{ or } \exists S \subseteq N^{<_\Gamma C} : S \models C$$

A first-order clause C is redundant w.r.t. the first-order clause set N (and $<_\Gamma$) iff

$$\forall C' \in \text{gnd}(C) : C' \text{ is redundant w.r.t. } \text{gnd}(N)$$

If redundancy does not hold, we call the corresponding clause non-redundant, or irredundant.

6.1 Learning Non-Redundant Clauses

First, we show that each learned clause is non-redundant w.r.t. the current clause set and induced ordering.

The most important consequence of this theorem that checking the learned clauses for redundancy criterions which are independent of the concrete induced orderings can be spared.

Such admissible criterions include subsumption, subsumption resolution and tautologies, as it is shown in the next subsection.

Theorem 6.2 (Non-redundant Clause Learning) *Let Γ denote the trail at a conflict in a regular run, $<_\Gamma$ the induced ordering, and assume the clause C is learned via the Backjump rule, and let N and U be the starting clause set and the set of learned clauses before the conflict, respectively.*

Then, C is not redundant w.r.t. $N \cup U$ and $<_\Gamma$.

Proof: Assume the first and last state in conflict resolution is

$$(\Gamma; N; U; k; (C_0; \sigma_0; \pi_0)) \Rightarrow^* (\Gamma'; N; U; k; (C; \sigma_1; \pi_1))$$

By soundness, $N \cup U \models C$ and each $C' \in \text{gnd}(C\sigma_1; \pi_1)$ is false w.r.t. both Γ' and Γ .

Now let $C' \in \text{gnd}(C\sigma_1; \pi_1)$ and assume there is an $S \subset \text{gnd}(N \cup U)$ such that $S \models C'$ and $S <_\Gamma C'$. Because of $S <_\Gamma C'$, each $C'' \in S$ has a defined truth-value w.r.t. Γ . If all $C'' \in S$ is true, then, by $S \models C'$, so is C' , a contradiction.

Thus, let $C'' \in S$ arbitrary such that C'' is false under Γ . We distinguish two cases whether Γ' is a strict subset of Γ , or equal to it.

First, if $\Gamma' \neq \Gamma$, at least one *Skip* had to be used, and C' contains no literal covered by the right-most literal of Γ . Neither does C'' , since $C'' <_\Gamma C'$. But then, C'' has a defined truth-value and it can only be true, as otherwise an earlier conflict detection would have been possible. A contradiction.

Second, assume $\Gamma' = \Gamma$. If the right-most literal is a decision, no false clause from $\text{gnd}(C_0\sigma_0; \pi_0)$ blocks this decision, and *Factorize* had to be applied several times followed by an application of case (2) of *Backjump*. (See also Lemma 5.5 on immediate conflicts.)

Let now C' such that it contains only a single top-level literal. Since case (2) of *Backjump* was used, such a clause from $\text{gnd}(C\sigma_1; \pi_1)$ exists. Since C'' is false and it was undefined before, it contains some top-level literals.

Since it was not a subject of *Propagate* before the right-most decision, it has to contain at least two such literals. But C' contains only one, and therefore $\text{def}(C') <_p \text{def}(C'')$ and $C' <_\Gamma C''$ must hold, a contradiction.

Finally, if $\Gamma = \Gamma'$ and the right-most literal is not a decision, the last rule had to be *Backjump* (case 1 or 2), and the same argumentation holds: If an assertive clause is learned, let C' an instance from $\text{gnd}(C\sigma_1; \pi_1)$ such that it contains only a single top-level literal. However, C'' must contain at least two top-level literals, which again leads to $C' <_\Gamma C''$, a contradiction. If $C = \perp$ is learned, it is smaller than any non-empty clause, and due to regularity, \perp is a newly learned clause. **Qed.**

6.2 Admissible Redundancies

Next, we show that the classic redundancy criterions *tautology*, *strict subsumption*, and *subsumption resolution* are *admissible redundancies in NRCL*, i.e. the clauses these rules remove are indeed redundant w.r.t. any induced ordering.

Proposition 6.3 (Tautology) *Let C a clause and N an arbitrary clause set.*

If $\models C$ holds, then C is redundant w.r.t. N .

Proof: Clearly, any ground instance of C is a ground tautology and redundant, since it follows from the empty set which "contains" only smaller clauses. **Qed.**

Furthermore, we also note that removing C has no effect on any run of the calculus, since no instance of C can be ever a conflict clause or imply an assignment.

Proposition 6.4 (Strict Subsumption) *Let C, D be clauses, σ a substitution, and N a set of clauses.*

If $C\sigma \subset D$, then D is redundant w.r.t. $N \cup \{C\}$.

Proof: Let $D\delta$ be a ground instance of D . Then $C\sigma\delta \subset D\delta$ and $C\sigma\delta <_{\Gamma} D\delta$ holds, for any induced ordering $<_{\Gamma}$. The latter holds, because $\text{def}(C\sigma\delta) <_p \text{def}(D\delta)$ holds in the abstract ordering.

Thus, $D\delta$ is redundant w.r.t. $\{C\sigma\delta\}$, and so is D w.r.t. $N \cup \{C\}$, and strict subsumption is admissible. **Qed.**

Similarly to tautology, removing a subsumed clause has little effect on the calculus, since whenever the subsumed clause is a conflict or a reason clause, the subsuming clause is either a conflict clause or implying the same assignment as well.

Proposition 6.5 (Subsumption Resolution) *Let C , D clauses, L a literal, σ a substitution, and N a clause set.*

If $C\sigma \subseteq D$ holds, then $D \vee \neg L\sigma$ is redundant w.r.t. $N \cup \{C \vee L, D\}$.

Proof: Redundancy clearly holds as D subsumes $D \vee \neg L\sigma$. Furthermore, we note that exchanging $D \vee \neg L\sigma$ with D in the presence of $C \vee L$ is a sound step. Thus, subsumption resolution as a rule for reducing a clause is admissible. **Qed.**

7 Termination and Completeness

Just as most related calculi, NRCL is a decision procedure for BS as well, under the regularity conditions of Definition 5.2. Below, we show that regular runs never get stuck and eventually terminate.

Proposition 7.1 *A regular run is never stuck, i.e. it terminates with the terminal rules, or one of the other rules is applicable.*

Proof: It is enough to show that, unless we already terminated, a rule is always applicable. First, we show that conflict search cannot get stuck.

If \perp is already in one of the clause sets, *Failure* is applicable and we terminate. Thus, w.l.o.g. assume $\perp \notin N \cup U$.

Assume Γ is total, i.e. defines each ground atom. Then I_{Γ} defines all ground atom occurring in $\text{gnd}(N)$, and it either satisfies N or there is a false ground clause from $\text{gnd}(N \cup U)$. In the first case, *Success* is applicable, and *Conflict* in the second case.

If Γ is not total, and some undefined ground literal is implied by some ground clause, *Propagate* is applicable. Otherwise, if no ground literal is implied and there is an undefined ground atom, we can always apply *Decide*. We note that decisions which define only a single ground atom are never blocked.

Second, assume we are resolving a conflict, i.e. the state indicator is $(C; \sigma; \pi)$ for some C , σ , and π . If the top literal in Γ is a decision and if $(C\sigma; \pi)$ is assertive, then *Backjump* is applicable. If it is not assertive, then either *Factorize*, or case (3) of *Backjump* is applicable.

If the top literal is a deduced literal, and neither does $C = \perp$ hold, nor is $(C\sigma; \pi)$ assertive - in these cases *Backjump* is applicable -, then we check the conditions of *Skip*. If *Skip* is not applicable, it satisfies the conditions of *Resolve*. Therefore, either *Skip*, *Factorize*, *Resolve* must be applicable in this case. **Qed.**

We show termination through a series of lemmas. First, we prove that both conflict search and conflict resolution always terminate:

Lemma 7.2 *Assume N , Σ and \mathcal{D} are all finite. Then, a conflict search phase of a regular run always terminates, i.e. leads either to a conflict or to termination.*

Proof: By the finiteness of Σ , we know that \mathcal{A}_Σ is also finite. Since a regular run is a series of sound steps, we also know that each application of *Propagate* and *Decide* defines at least one formerly undefined ground atom.

Thus, a regular run eventually exhausts these rules, and, since it cannot get stuck by Proposition 7.1, one of the rules *Failure*, *Success*, or *Conflict* has to be applied. And thereby, the conflict search phase in question ends. **Qed.**

Lemma 7.3 *Assume N , Σ and \mathcal{D} are all finite. Then, a conflict resolution phase of a regular run always terminates, i.e. leads to the application of Back-jump in finitely many steps.*

Proof: Let us assign to each intermediate state $(\Gamma; N; U; k; (C; \sigma; \pi))$ in a conflict resolution the tuple $(\#(\Gamma); \text{gnd}(C\sigma; \pi))$ as a measure, where $\#(\Gamma)$ denotes the number of elements in Γ .

Let us order these tuples with the lexicographical ordering $<_{lex}$ based on the canonical ordering over non-negative integers and $<_0$ where $<_0$ denotes both the ordering induced by the trail after finding the conflict, and its multiset extension. This ordering is well-founded.

We note that conflict resolution cannot get stuck, see Proposition 7.1. Therefore, it is enough to show that each application of the rules *Skip*, *Resolve*, and *Factorize* strictly decreases our measure.

Skip strictly decreases the size of Γ , and therefore our measure as well. In the case of *Resolve* and *Factorize*, it is enough to give a function satisfying the conditions of Proposition 2.26 between the false instances on the two sides, i.e. a function γ which assigns ground clauses from the right-hand side conflict-set to larger ground clauses from the left-hand conflict-set.

First, assume we apply *Resolve* to the state

$$(\Gamma, (L' \cdot \sigma'; \pi')^{C' \vee L'}; N; U; k; (C \vee L; \sigma; \pi))$$

and we get

$$(\Gamma, (L' \cdot \sigma'; \pi')^{C' \vee L'}; N; U; k; ((C \vee C')\eta_0; \sigma^*; \pi\eta, \pi'\eta))$$

where $\eta = \text{mgu}(L'\sigma', \neg L\sigma)$, $\eta_0 = \text{mgu}(L', \neg L)$, and σ^* such that $\eta_0\sigma^* = \sigma\sigma'\eta$. For the sake of readability, let us introduce the symbols $\alpha = \sigma\sigma'\eta$ and $\pi^* = \pi\eta, \pi'\eta$.

Now, let β be a grounding substitution such that $(C \vee C')\alpha\beta \in \text{gnd}((C \vee C')\alpha; \pi^*)$. Since it was derived via resolution, there is a corresponding valid ground resolution step with premises

- $C_1 \vee L_1 \in \text{gnd}((C \vee L)\alpha; \pi\eta)$
- $C_2 \vee L_2 \in \text{gnd}((C' \vee L')\alpha; \pi'\eta)$

where we assume L_1 and L_2 are the literals corresponding to L and L' , respectively. Since we apply resolution, we also know that $L_1 = \neg L_2$, and $(C \vee C')\alpha\beta = C_1 \vee C_2$.

By the definition of sound states and *Propagate*, we know that C_2 contains only literals which were defined before the last assignment, and thus, $C_2 <_0$

$(\neg)L_2$, and therefore $C_2 <_0 L_1$. Then, $C_1 \vee C_2 <_0 C_1 \vee L_1$ must hold, and thus, we shall define $\gamma(C_1 \vee C_2)$ as $C_1 \vee L_1$.

Since γ can be defined over the whole $\text{gnd}((C \vee C')\alpha; \pi^*)$ and $\text{gnd}((C \vee L)\alpha; \pi\eta)$ is a subset of $\text{gnd}((C \vee L)\sigma; \pi)$, we can apply Proposition 2.26, and we get

$$\text{gnd}((C \vee L)\sigma; \pi) >_0 \text{gnd}((C \vee C')\alpha; \pi^*)$$

and our measure strictly decreases, as the size of the trail is unchanged. The proof for *Factorize* is analogous. **Qed.**

Next, we show that only finitely many new clauses can be learned thanks to our non-redundancy results in Theorem 6.2.

Lemma 7.4 *If N , Σ and \mathcal{D} are finite, a regular run can only learn finitely many new clauses.*

Proof: We use Higman's Lemma [17] to prove this claim. The lemma states that given an infinite sequence w_1, w_2, \dots of words over a finite alphabet, there is always an index i and a subsequent index j such that the word w_i is *embedded* into w_j , i.e. after deleting some letters from w_j we can get w_i .

Now, consider \mathcal{A}_Σ . Since Σ and \mathcal{D} are finite, both the set of ground atoms and ground literals over Σ and \mathcal{D} are finite. The latter serves as the finite alphabet for our proof.

Since every learned clause is non-redundant at the time they are learned, by Theorem 6.2, we can assign a non-redundant ground instance to any learned clause, by the definition of redundancy.

Assume we learn infinitely many clauses, and let us consider the assigned ground clauses C_1, C_2, \dots , where C_1 is assigned to the clause learned at the first conflict, C_2 to the clause learned at the second, and so on.

Now, take any term ordering $>$, order the literals of the clauses, and assign this ordered sequence of literals to each clause. Let us denote this word over the alphabet of ground literals by $w(C)$ for every ground clause C .

Then, by Higman's Lemma, there are indices $i < j$ such that $w(C_i)$ is embedded in $w(C_j)$. But it means that $C_i \subseteq C_j$, i.e. C_j is strictly subsumed by or equal to C_i .

The admissibility of strict subsumption was proven in Proposition 6.4, and clearly an already present ground clause cannot be non-redundant either, for any induced ordering. Thus, C_j cannot be redundant at the j th conflict, a contradiction. **Qed.**

Finally, we show termination, and state the main result as a corollary.

Theorem 7.5 (Termination) *A regular run always terminates if N , Σ and \mathcal{D} are finite.*

Proof: First, we note that a run can be seen as a series of conflict search and conflict resolution phases, which ideally ends with a terminal rule. By Lemma 7.2, Lemma 7.3, and Proposition 7.1, we know that each phase ends after finitely many steps without getting stuck.

Thus, an infinite run must be an infinite series of conflict search and resolution sequences. Since each conflict resolution ends with *Backjump*, it would imply that infinitely many new clauses are learned. But it contradicts Lemma 7.4. **Qed.**

Corollary 7.6 (Decision Procedure) *Regular runs provide a decision procedure for the Bernays-Schönfinkel fragment if N , Σ and \mathcal{D} are finite.*

I.e. every regular run terminates after finitely many steps with Failure, or Success, for an unsatisfiable, or satisfiable clause set N , respectively.

Proof: It follows from Proposition 7.1 and the Theorems 4.4 and 7.5. **Qed.**

8 Towards Implementation

This far we considered mostly our calculus in an abstract fashion, and it is enough to establish the results of the previous chapters.

Here, we elaborate some details regarding the constraints, and refine some steps to bring NRCL closer to practical application. In particular, we provide an abstract algorithm for exhaustive propagation, to highlight some important difficulties and expensive steps in the calculus.

However, this section does not aim to provide a complete abstract algorithm for regular runs, we only briefly address some challenges and propose some solutions and approaches, which provides us a starting point for later implementation and experimentation.

8.1 Free Variables

The definition of normal form for constrained literals demands the left-hand side of a constraint to contain only variables occurring in the constrained literal. Our calculus derives new assignments, i.e. new constrained literals for Γ , by applying resolution between the literals in Γ and the clauses in $N \cup U$.

However, even after normalization, the resulting candidate $(L \cdot \sigma; \pi)^C$ might contain free left-hand side variables, i.e. variables which occur in the reason clause instance $C\sigma$, and still occur in $\text{lvar}(\pi)$, but do not occur in $L\sigma$. The following example demonstrates this behavior.

Example 8.1 *Let us take*

$$N = \{C_1 : \neg Q(x, x), C_2 : \neg Q(x, y) \vee \neg Q(x, z) \vee P(y, z)\}$$

And assume that after an application of Propagate and Decide we get the trail

$$\Gamma = (\neg Q(x, x); \top)^{C_1}, (Q(x, y); (x, y) \neq (v, v))^1$$

Now, applying Propagate between Γ and the clause C_2 , we get the constrained literal

$$(P(y, z); (x, y) \neq (v, v) \wedge (x, z) \neq (w, w))$$

Over $\mathcal{D} = \{a, b\}$, this constraint is satisfiable, the cover-set is $\{P(a, a), P(b, b)\}$, and after eliminating the free variable x we get the constrained literals

$$(P(y, z); y \neq a \wedge z \neq a) \text{ and } (P(y, z); y \neq b \wedge z \neq b)$$

Semantically, these variables are to be treated as existential variables, of course. These variables cause two problems.

First, in the presence of these existentially handled variables our constrained literal set for difference defined in Lemma 2.16 is no longer valid. In particular disjointness is no longer guaranteed.

A simple way to overcome this issue is to split the resulting literal into a set of literals by instantiating the free left-hand side variables with relevant constants, as seen in the example above. This elimination procedure results in a set of not necessarily disjoint constrained literals.

Second, while eliminating these variables is a solution, we still need to store the instantiating assignments. This information is used when applying the rules *Resolve* and *Factorize* during conflict resolution. This is already accomplished through using *closures* as introduced in Section 3.

8.2 Indexing Scheme

In the propositional setting, the watched literal scheme watches two literals in every non-unit clauses. These literals are assumed to be true or undefined under the current model assumption, or all literals but a single watched literal are false in the clause.

Whenever a new assignment makes a watched literal false, we attempt to find a new non-false literal. If it is not possible, the other watched literal is propagated resulting either in a new assignment or a new conflict clause.

This scheme enables efficient propagation at small computational costs as it cuts back the number of clauses we have to consider after a new assignment and requires no additional bookkeeping during backtrack.

When lifting the scheme, we have to keep in mind that manipulating our constraints is more expensive. Therefore, a direct lifting of the technique by exactly maintaining which literals are watched in the different instances of a clause would be too expensive for our purposes.

Here, we propose a lightweight approach which uses two levels of indexing the literals of the current clause set. Every clause is indexed by one of these levels, but not both.

The first level attempts to mimic the two-watched-literal scheme, and indexes only two literals in the clauses. We can choose the interpretation of watching a literal L as an approximation of *cannot be false* by selecting one of the following:

- $\nexists (L' \cdot \sigma; \pi)^\alpha \in \Gamma : \exists \text{mgu}(\neg L' \sigma, L)$
- $\nexists (L' \cdot \sigma; \pi)^\alpha \in \Gamma : \exists \delta = \text{mgu}(\neg L' \sigma, L)$ and $\pi \delta \neq \perp$
- $\nexists (L' \cdot \sigma; \pi)^\alpha \in \Gamma : \exists \delta = \text{mgu}(\neg L' \sigma, L)$ and $\pi \delta$ is not empty

Obviously, the last choice is the most expensive and the first two should be preferred.

Whenever a new assignment is made, we first try to adjust the watched literals on level one. If a clause contains no longer two appropriate literals, we push it to the second level. On this level we index all literals of the clauses, e.g. in a context tree with top-level symbol hashing.

Putting clauses back to level one can be done either by maintaining an activity heuristics and time to time manually check for watchable literals, or managing lists of pointers for all clause-literals to relevant assignments on the trail.

This topology should make propagation cheaper, and in particular using level one should make it easier to ignore clauses irrelevant w.r.t. the recent assignments.

8.3 Finding Candidates

Before we propose an abstract algorithm for exhaustive propagation, we introduce a simple derivation system for finding candidates. Of course, in the actual implementation this system will be replaced by more efficient algorithms on the indexing structures.

The rules work on tuples of the form $(C; \sigma; \pi)^i$ where

- C is a clause, a subclause of some initial clause C_0 from the current clause set
- σ is a substitution over $\text{var}(C_0)$
- π is a dismatching constraint
- i is the number of application of the last assignment of the trail, which has relevance in the next section

The initial tuple for a clause $C_0 \in \mathbf{N} \cup \mathbf{U}$ is $(C_0; \emptyset; \top)^0$ and we try to resolve each literal in C with the following rule:

$$(C \vee L; \sigma; \pi)^i \Rightarrow_{\Gamma} (C; \sigma\theta; \pi\theta \wedge \pi'\theta)^{i'}$$

Where there is a $(\neg L' \cdot \sigma'; \pi')^\alpha \in \Gamma$ such that

- $\exists \theta = \text{mgu}(L\sigma, L'\sigma')$
- $(\pi\theta \wedge \pi'\theta) \neq \perp$ and normalized
- i' is $i + 1$ if $(\neg L' \cdot \sigma'; \pi')^\alpha$ is the last assignment in Γ , and i otherwise

Applying this rule we can get candidates for the rules *Conflict*, and *Propagate* by deriving respectively tuples of the form

- $(\perp; \sigma; \pi)^i$, or
- $(L; \sigma; \pi)^i$

We note that non-emptiness is not checked fully, only a cheaper precondition of it. Free left-hand side variables and already defined instances are not removed either.

8.4 Exhaustive Propagation

In this section, we propose the abstract algorithm **PROP** for exhaustive propagation with conflict detection. It basically processes a queue **PQ** of candidates for new assignments. As an invariant, we assume each constrained literal in the queue

1. has a normalized non- \perp constraint
2. consistent with the current Γ
3. contains no free left-hand side variable

PROP

Initially, this queue consists of the literals induced by the unit clauses. Unit clauses has to be checked for contradiction prior calling **PROP**. When calling after decisions, **PQ** is assumed to contain the immediate consequences of the decision. Checking for blocking should generate this set anyway.

PROP processes the literals on **PQ**. First, it removes already defined instances by calling the function **DIFF**. This produces a set of disjoint and undefined constrained literals, each of which is a valid subject of *Propagate*. See Section 2.3 for the definition of the difference operation “-”, and see below the abstract algorithm for **DIFF**.

These literals are then checked for emptiness, added to Γ and set to true. Their consequences - conflicts and new candidates for **PQ** - are then generated by **addConsequences**.

We continue this process until **PQ** gets empty, or a conflict is found. The first indicates the finished exhaustive application of *Propagate*, and *Decide* can be called. In this case we return *true*. And in the latter case, we return *false*, and the found conflict is stored in **conflictSet**.

On the course of this section, we might use the symbol ℓ to denote annotated constrained literals, and the following auxiliary functions:

- **pop**: removes an element of a queue, list, or set
- **notEmpty**: carries out a full non-emptiness check for a constraint or constrained literal
- **addAssignment**: adds a new assignment to Γ (and its indexing structures)
- **cUNIF**(ℓ , Γ): Finds the literals in Γ which are unifiable with ℓ , and returns an array of them and its size
- **NF**: normalizes a constraint, constrained literal, or a set of constrained literals, as described in Subsection 2.2. In the latter case, it removes resulting literals with \perp -constraints.
- **freeLVars**: produces the set of free left-hand side variables of a constrained literal
- **selectOne**: randomly, or heuristically selects an element of a set, or a list
- **adjustLevel1**: adjusts the first index level for clauses after a new assignment given as parameter, as described in Section 6.2.
- **getCandidates**: provides the list of indexed clauses which contains a literal unifiable with the complement of a given literal

DIFF

It iteratively removes the already defined instances from the proposed assignment. The result is a set of disjoint and undefined constrained literals with non- \perp constraints.

Function PROP(N, U, Γ, PQ)

```

1 while  $PQ \neq \emptyset$  do
2    $\ell = (L \cdot \sigma; \pi)^C \leftarrow \text{pop}(PQ)$ ;
3    $\Delta \leftarrow \text{DIFF}(\ell, \Gamma)$ ;
4   foreach  $\ell' = (L' \cdot \sigma'; \pi')^C \in \Delta$  do
5     if  $\text{notEmpty}(\ell')$  then
6       // Applying Propagate:
7        $\text{addAssignment}(\Gamma, \ell')$ ;
8       if  $\text{addConsequences}(N, U, \Gamma, \ell', PQ) = \text{false}$  then return false;
9     end
10  end
11 return true;

```

Function DIFF(ℓ^*, Γ)

```

1  $(\vec{\ell}, k) \leftarrow \text{cUNIF}(\ell^*, \Gamma)$ ;
2  $\Delta_0 \leftarrow \{\ell^*\}$ ;
3 for  $i = 1, \dots, k$  do
4    $\Delta_i \leftarrow \text{NF}(\Delta_{i-1} - \vec{\ell}[i])$ ;
5 end
6 return  $\Delta_k$ ;

```

elimFV

An auxiliary function for finding new candidates. It iteratively removes the free left-hand side variables, and only keeps the literals with non- \perp normalized constraints.

addConsequences

Finally, **addConsequences** checks whether a new assignment produces a conflict and generates new candidates for PQ . It returns *true* if no conflict is found, and *false* otherwise. If a conflict is found, it is saved in **conflictSet**.

We distinguish two types of conflicts. It is easy to see, that if the new assignment is used only once in deriving a conflict, then PQ must already hold an unprocessed candidate which is falsified by the new assignment. Thus, we check PQ first for a contradiction, and start generating new candidates with \Rightarrow_Γ only afterwards.

We then use the derivation system of 6.3 to derive new constrained literals. We only consider derivations where the latest assignment has to be used at least once. If it is used only once we can be sure the new literal is not false. If it is not the case, we check for a possible conflict.

As stated before, in the actual implementation the proper retrieval algorithms will eliminate the inefficiency of considering all derivations.

Finally, the new candidates are tested for free variables, and they are removed if there are any.

Function elimFV(ℓ_0)

```

// Prereq:  NF( $\pi_{\ell_0}$ )  $\neq \perp$ 
1  $\Delta, \Delta^* \leftarrow \{\ell_0\}, \emptyset;$ 
2 while  $\Delta \neq \emptyset$  do
3    $\ell = (L \cdot \sigma; \pi)^C \leftarrow \text{pop}(\Delta);$ 
4   if freeLVars( $\ell$ ) =  $\emptyset$  then
5      $\Delta^* \leftarrow \Delta^* \cup \{\ell\};$ 
6   end
7   else
8      $x \leftarrow \text{selectOne}(\text{freeLVars}(\ell));$ 
9     // Instantiation with each constant from the domain:
10    foreach  $d \in \mathcal{D}$  do
11       $\pi' \leftarrow \text{NF}(\pi\{x \leftarrow d\});$ 
12      if  $\pi' \neq \perp$  then  $\Delta \leftarrow \Delta \cup \{(L \cdot \sigma\{x \leftarrow d\}; \pi')^C\};$ 
13    end
14  end
15 return  $\Delta^*$ 

```

8.5 Picking the Next Decision

When making a new decision, we pick a candidate $(L; \pi)$, remove all the already defined instances, and then test all immediate conflicts for blocking.

If there is a blocking conflict, we might then either pick an entirely new decision candidate, or try to *fix* $(L; \pi)$ by instantiating some variables in L , and thereby generating a new set of candidates.

This can be achieved by picking a blocking ground instance $C' = C\delta$ which contains $\neg L\delta_1, \neg L\delta_2$ such that both $L\delta_1, L\delta_2 \in \text{gnd}(L; \pi)$ and $L\delta_1 \neq L\delta_2$ holds. Now, choose a variable for which $x\delta_1 \neq x\delta_2$, and split $(L; \pi)$ into $(L\{x \leftarrow x\delta_1\}; \pi\{x \leftarrow x\delta_1\})$ and $(L; \pi \wedge x \neq x\delta_1)$. By instantiating further variables, we eventually get a decision which is not blocking, since a ground decision is always suitable.

A non-blocking decision is then added to Γ , and whether we found a non-blocking conflict or not, we continue with conflict resolution or with calling **PROP** after generating the immediate propagation candidates in a similar way as in **addConsequences**.

Initially, the set of decision candidates are generated from the literals occurring in N . This set can be later refined by the above steps, and individual candidates might be substituted with sets of new candidates.

Since removing defined instances is always relative to the current Γ , it has to be guaranteed that the set of all possible candidates covers the original set. It can be ensured for example by keeping a trail for these refinement steps as well, and re-roll them in parallel with the backtracking procedure.

8.6 Ranking Literals

Most current SAT solvers also employ variable selection schemes based on dynamic ranking of propositional variables. This technique rewards variables in-

Function addConsequences(N, U, Γ, ℓ', PQ)

```

1   $(L' \cdot \sigma'; \pi')^\alpha \leftarrow \ell'$ ;
   // Step 1: Check Type-1 Conflicts
2  foreach  $\ell = (L \cdot \sigma; \pi)^C \in PQ$  do
3  |   if  $\exists \delta = \text{mgu}(\neg L\sigma, L'\sigma')$  and  $\text{NF}(\pi\delta \wedge \pi'\delta) \neq \perp$  and
   |   notEmpty $(\pi\delta \wedge \pi'\delta)$  then
4  | |    $\text{conflictSet} \leftarrow (C; \sigma\delta; \pi\delta \wedge \pi'\delta)$ ;
5  | |   return false;
6  |   end
7  end
   // Step 2: Adjust indexing level 1
8   $\text{adjustLevel1}(N, U, \ell')$ ;
   // Step 3: Generate consequences
9  foreach  $C \in \text{getCandidates}(N \cup U, \ell')$  do
10 |   foreach  $\text{derivation } (C; \emptyset; \top)^0 \Rightarrow_\Gamma^* (L^*; \sigma^*; \pi^*)^i$  with  $i \geq 1$  do
   |   // Check Type-2 Conflicts
11 |   if  $i \geq 2$  then
12 | |   if  $\exists \ell = (L \cdot \sigma; \pi)^\beta \in \Gamma$  such that  $\exists \delta = \text{mgu}(\neg L\sigma, L^*\sigma^*)$  and
   | |    $\text{NF}(\pi\delta \wedge \pi^*\delta) \neq \perp$  and notEmpty $(\pi\delta \wedge \pi^*\delta)$  then
13 | | |    $\text{conflictSet} \leftarrow (C; \sigma^*\delta; \pi\delta \wedge \pi^*\delta)$ ;
14 | | |   return false;
15 | |   end
16 |   end
17 |    $\ell^* \leftarrow (L^* \cdot \sigma^*; \pi^*)^C$ ;
18 |   if  $\text{freeLVars}(\ell^*) = \emptyset$  then  $PQ \leftarrow PQ, \ell^*$ ;
19 |   else  $PQ \leftarrow PQ, \text{elimFV}(\ell^*)$ ;
20 |   end
21 end
22 return true;

```

volved in recent conflicts, and proved itself efficient in the propositional context.

Following the footsteps of the now classic *decaying variable sum*, we reward the literals involved in the clause learning phase following the latest conflict.

This is accomplished by maintaining a list of literals and scores. Whenever some literal L is added to the clause of the intermediate state, we add a pair $(L; v)$ to this list.

To focus on recent conflicts, we increase v gradually, and occasionally we reset v to some initial value and normalize the list. The latter can be triggered upon reaching some extreme value, automatically after a certain number conflicts, or at restarts.

Restarts are commonly used in SAT solvers to redirect the focus of the search using the learned clauses and the current variable scores. Applying it only finitely many times does not violate completeness.

Then, whenever we need to choose a new decision, we rank the candidates by combining the scores belonging to literals which are unifiable with the candidate in question. As an example we propose addition or maximum. We then choose the literal with the highest combined score.

8.7 Clause Learning and Backjumping

As the conflicts are now discovered, every conflict-set uniquely assigns a Γ -assignment to each literal of the conflict clause. This makes detecting assertiveness easy and spares us a number of emptiness checks, as they are already done during conflict detection. This way, the only non-deterministic choice is the application of *Factorize* versus *Resolve*, when both is applicable.

Once a new clause is learned, a suitable backtrack level is needed. Should we learn only the ground clauses in $\text{gnd}(C\sigma; \pi)$ when the last conflict-set is $(C; \sigma; \pi)$, we could determine the backtrack position at ease, similarly to the propositional solvers.

But we learn the more general C , and the right backtrack position has to be computed from all the instances of C . We have to consider all conflicting instance of C w.r.t. Γ , and for each instance, we have to determine a minimal backtrack position. Then, we backjump to the minimum of these positions.

Without providing more details, we only note that some instances produce new assignments after backtrack, some might block existing decisions, and some might even be new conflicts after backjump.

9 Related Work

In this section, we briefly compare NRCL to existing solutions. As Bernays-Schönfinkel problems can be successfully handled with finite model finders as well, we cover both BS-specific techniques and more general finite model building approaches. In the case of the latter systems, we focus on their behavior on the Bernays-Schönfinkel fragment.

The first successful approaches to finite model building were *Mace* and *SEM*, see e.g. [39]. The early version of *Mace* flattens and grounds the given clause set, and passes it on to a CDCL-based SAT solver. This approach is developed further by *Paradox* [10].

Compared to this approach, we work directly with the first-order clause set instead of the often exponentially larger set of ground instances.

The latest version of *Mace* [26] follows the approach of *SEM* [42] and *FINDER* [37]. Instead of generating the ground instances, it maintains the function and predicate tables, and fills them out using a sophisticated backtracking algorithm.

Compared to this approach, we represent the model implicitly via constrained literals, and let the learned clauses guide our calculus.

Over the last decade several attempts were made to lift CDCL and its ancestor, DPLL - a calculus using backtracking instead of backjumping and clause learning. *Model Evolution* [7] and its implementation *Darwin* [5] represents a model with a set of first-order literals, called *context*, and detects conflicts using syntactic concepts weaker than the full-fledged semantics based on the induced interpretation. This potentially leads to longer derivation before detecting a false clause.

It is refutationally complete over first-order clauses and provides a decision procedure for the Bernays-Schönfinkel fragment. Its extension [6] enriches the calculus with learning lemmas at conflicts, and uses backjumping instead of the original backtracking approach.

Compared to *Model Evolution*, NRCL relies on the full-fledged semantics, and we learn only non-redundant clauses. It is not clear if the latter holds for *Model Evolution*, especially the admissibility of the classic criteria needs in-depth considerations.

Finally, it was shown in [14] that using contexts might result in exponentially larger model representations. We note that this result holds for the general case with function symbols, but in our setting e.g. the constrained literal

$$(P(x_1, x_2, \dots, x_k); x_1 \neq x_2 \wedge x_2 \neq x_3 \wedge \dots \wedge x_{k-1} \neq x_k)$$

whose size is $O(k)$, requires a representation of size at least $O(k^2)$ as a context. Thus, at least a quadratic relation holds even for the Bernays-Schönfinkel fragment.

DPLL(SX) [34] attempts to lift CDCL to BS in the same manner as we do, has an almost identical rule set, and uses substitution sets represented by BDDs as constraints. Substitution sets provide an explicit way to represent models.

It is well-known that in the general setting with function symbols implicit representations have stronger expressive power [33][24]. In our setting, explicit representations have the potential to be exponentially larger than the corresponding implicit representations.

The following simple example demonstrates this claim. Over $\mathcal{D} = \{a, b, c\}$, consider the constrained literal

$$(P(x_1, x_2, \dots, x_k); x_1 \neq x_2 \wedge x_2 \neq x_3 \wedge \dots \wedge x_{k-1} \neq x_k)$$

Then it is easy to see that the corresponding explicit representation is made up of all the ground instances covered by this literal.

Therefore, while the size of the implicit representation increases linearly in k , the size of the corresponding explicit representation is $O(2^k)$, i.e. increases exponentially in k .

The authors of this paper are convinced that this exponential blow-up happens whenever the implicit representation has no finite explicit representation

(see [33][24] for details) in the language enriched with a function symbol. However, this conjecture needs further consideration, and we leave it for future work.

Furthermore, compared to $DPLL(SX)$ our approach is more modular as it allows the use of an arbitrary constraint language, restricted only by the operations we expect. Dismatching constraints can be extended beyond the Bernays-Schönfinkel fragment easily, while in the case of BDD-encodings, it is not trivial.

$DPLL(SX)$ also lacks the concept for blocking, and applies an explicit refine rule instead. As a side effect, it learns nothing from conflicts which lead to blocking clauses, and in these cases it abandons conflict resolution and refines the last decision. Finally, we also address redundancy, and exploit the non-redundancy result to show termination, which we consider a valuable addition.

The most recent calculus $SGGS$, introduced in [9], promises a semantically guided, goal sensitive, model-based proof system. It uses simple constraints, so-called *standard forms*, conjunctions of negative atomic constraints of the form $x \neq y$, or $\text{top}(x) \neq f$.

Then, a model is represented by a sequence of constrained clauses with selected literals. This sequence overrides a given initial interpretation I , which serves both as initial model assumption and as semantic guidance for the calculus.

The procedure then keeps expanding this sequence in order to satisfy more and more clauses, and handles contradictions via resolution and splitting the constrained clauses to maintain an invariant - every literal in every clause in the sequence must have *either only false, or only true instances* w.r.t. I and the constraints.

NRCL utilizes a more expressive constraint language, which allows tuples to be used. This results in less fragmentation of the representation, i.e. $SGGS$ might need several constraints in standard form to express a single dismatching constraint of our calculus.

This allows us to learn more general clauses, and also potentially decreases the size of the representation. Our model representation relies on constrained literals instead of clauses, and we consider it to be more explicit than the approach of $SGGS$ which requires identifying the constrained instances of the clauses which are indeed producing new assignments.

Finally, the resolution applied by $SGGS$ only repairs the model, it can be discarded later as the search progresses, and the splittings applied to maintain the invariants also forces the result of resolution to be more specific, more local. Compared to this, our calculus learns and saves new clauses, uses backjumping, and we proved these clauses are non-redundant.

We also mention *geometric resolution* [29] which uses a special normal form called *geometric normal form*. In this calculus the formulas themselves constitute the rules of a system based on backtracking. Through the inference *geometric resolution* it also provides a way to learn new formulas. The transformation to geometric normal form also includes flattening, which our approach avoids.

The calculus *Inst-Gen* [16] and its implementation *iProver* [21] has been quite successful at solving Bernays-Schönfinkel problems, and competitive even for the first-order fragment. It generates a propositional approximation of the clause set by instantiating all the variables with constants, and passes it on to

a CDCL-based SAT solver.

Unsatisfiability of the approximation entails the unsatisfiability of the original problem. On the other hand, if an abstract model is generated, it is used to guide the calculus to add proper instances of the original clauses, which refines the propositional abstraction.

This procedure is continued then, until either unsatisfiability is proven, or saturation is achieved, which implies that the abstract model can be lifted to a first-order model for the original clause set.

The algorithm is further enhanced by using dismatching constraints, and applying redundancy elimination based on generating first-order resolvents for subsumption with a theorem prover, and finding simplification candidates efficiently with ground reasoning.

Compared to *iProver*, our approach is fine-grained, as the evaluation and refinement of our abstraction happen interleaved with the other reasoning steps. Furthermore, we work directly with the original clause set, and our trail always corresponds to a consistent first-order model candidate.

In addition to the theoretical comparison, we also ran a small experiment for models represented by literals of the form $(P(x_1, \dots, x_k); x_1 \neq x_2, \dots, x_{k-1} \neq x_k)$. The clause set

$$\begin{aligned} &Q(x, x), \neg Q(a_1, a_2), \neg Q(a_2, a_3), \dots, \neg Q(a_{n-1}, a_n), \\ &\neg P(x_1, x_1, x_3, \dots, x_k), \dots, \neg P(x_1, x_2, \dots, x_{k-1}, x_{k-1}), \\ &\neg Q(x, z) \vee Q(x, y) \vee Q(y, z), P(x_1, \dots, x_k) \vee Q(x_1, x_2) \vee \dots \vee Q(x_{k-1}, x_k) \end{aligned}$$

has a model where the positive atoms are represented by the constrained literals $(P(x_1, \dots, x_k); x_1 \neq x_2, \dots, x_{k-1} \neq x_k)$ and $(Q(x, x); \top)$. NRCL directly finds this model, i.e. without backjumping even once, by exhaustively applying propagation, making a single decision on $P(x_1, \dots, x_k)$ and finally setting all undefined $Q(x, y)$ literals to false. Furthermore, any regular run would find a similar model without backjumping even once.

We tested this clause set with the available state-of-the-art provers *Darwin* (1.4.5) and *iProver* (0.8.1). The experiments were carried out on a Debian Linux (4.7.2-5) Intel (Xeon E5-2680, 2.7GHZ) computer with 256GB physical memory. For $n = 7$ and $k = 5, 7, 9$, *Darwin* needs 0.2, 8.1, 518 seconds to find a model, respectively. For $k = 7$ and $n = 7, 10, 13$, *Darwin* needs 8.1, 62, 347 seconds to find a model, respectively. For $k = 9$ and $n = 7, 10, 13, 16, 19, 22$, *iProver* needs 0, 0.2, 21, 39, 116, 718 seconds to find a model, respectively.

In the case of *Darwin*, these results show an exponentially growing solution time w.r.t. k (the arity of P) or n (the domain size). *iProver* is robust against increasing k but not against increasing n , where it also shows an exponential growth. This shows that our model representation is not subsumed by either *Darwin* or *iProver*.

Finally, even general purpose first-order theorem provers implement specialized techniques to handle Bernays-Schönfinkel problems.

Generalisation introduced in [31] for *Vampire* is an additional technique for resolution-based saturation. It infers $P(x)$ if $P(c)$ has been established for all relevant constant c . Coupled with efficient sort inference, it has the potential to exponentially speed up theorem proving.

The technique introduced in [18] for *SPASS* employs a combination of restricted superposition on Horn clauses, and *labelled splitting* [15] on non-Horn clauses.

Compared to these approaches, NRCL maintains a model candidate, it is restricted to learn clauses only at conflicts and only non-redundant ones, does not rely on Horn clauses, and the implicit branchings through decisions and backjumps are more elaborate and guided by the model search, compared to the splitting techniques employed by first-order theorem provers. However, we note that for some problem classes finite superposition saturation is still superior to explicit model generation, see e.g. superposition for knowledge bases in [38].

10 Conclusion

In this paper, we proposed the decision procedure NRCL for the Bernays-Schönfinkel fragment. Our approach represents a model candidate as a set of constrained literals, and derives a model or a proof of unsatisfiability through a series of decisions, propagations, and learning new clauses.

Our work closely relates to $DPLL(SX)$ [34], which introduces a similar calculus, and the more recent calculus $SGGS$ [9]. Compared to earlier work in this direction, we investigated the standard redundancy notion w.r.t. the ordering induced by the current trail.

One of the main contributions of NRCL over existing work is that, by design, we can prove our learned clauses to be non-redundant, i.e., any learned clause makes progress towards finding a model or a refutation, because it eliminates at least one potential model. In general, we consider this a key property for automated reasoning calculi.

Projecting NRCL down to propositional logic proves this property for CDCL with respect to our notion of redundancy. Our notion also admits techniques like subsumption and subsumption resolution, which are important in both SAT solving and first-order theorem proving. We see this as a strong indication that a future implementation will also contribute to the state of the art.

In Section 8, we addressed some of the difficulties of this approach, and provided details for implementation. Finally, we gave a brief comparison to the existing solutions in Section 9.

As future research, the immediate goal is to make an efficient implementation of NRCL. This includes developing suitable and efficient term indexing structures, possibly revising the constraint language, and defining concrete and efficient heuristics for selecting decisions.

On the other hand, the long-term goal of our research is to extend this calculus beyond Bernays-Schönfinkel. The next step into this direction is to enrich our calculus with function symbols and sorts to handle the *non-cyclic fragment* introduced in [23]. This class still has the finite Herbrand model property, thus, our results will directly extend to this fragment.

The further goals are to consider other decidable fragments, to introduce equality into our calculus, and finally to extend our work to finite model finding.

References

- [1] A. Armando, P. Baumgartner, and G. Dowek, eds. *Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia*,

August 12-15, 2008, *Proceedings*, 2008, *LNCs* 5195. Springer.

- [2] L. Bachmair and H. Ganzinger. Resolution Theorem Proving. In Robinson and Voronkov [35], pp. 19–99.
- [3] L. Bachmair, H. Ganzinger, C. Lynch, and W. Snyder. Basic Paramodulation and Superposition. In D. Kapur, ed., *Automated Deduction - CADE-11, 11th International Conference on Automated Deduction, Saratoga Springs, NY, USA, June 15-18, 1992, Proceedings*, 1992, *LNCs* 607, pp. 462–476. Springer.
- [4] L. Bachmair, H. Ganzinger, C. Lynch, and W. Snyder. Basic Paramodulation. *Inf. Comput.*, 121(2):172–192, 1995.
- [5] P. Baumgartner, A. Fuchs, and C. Tinelli. Darwin: A Theorem Prover for the Model Evolution Calculus. In S. Schulz, T. Tammet, and G. Sutcliffe, eds., *Proceedings of the 1st Workshop on Empirically Successful First Order Reasoning (ESFOR’04)*, 2004, IJCAR 2004 Workshop Proceedings, pp. 1–24. UCC.
- [6] P. Baumgartner, A. Fuchs, and C. Tinelli. Lemma Learning in the Model Evolution Calculus. In M. Hermann and A. Voronkov, eds., *LPAR*, 2006, *LNCs* 4246, pp. 572–586. Springer.
- [7] P. Baumgartner and C. Tinelli. The Model Evolution Calculus. In F. Baader, ed., *CADE*, 2003, *LNCs* 2741, pp. 350–364. Springer.
- [8] A. Biere, M. Heule, H. van Maaren, and T. Walsh, eds. *Handbook of Satisfiability*, 2009, *Frontiers in Artificial Intelligence and Applications*, vol. 185. IOS Press.
- [9] M. P. Bonacina and D. A. Plaisted. SGGS theorem proving: an exposition. *Notes of the Fourth Workshop on Practical Aspects in Automated Reasoning (PAAR), Seventh International Joint Conference on Automated Reasoning (IJCAR) and Sixth Federated Logic Conference (FLoC), Vienna, Austria, July 2014.*, 2014.
- [10] K. Claessen and N. Srensson. New Techniques that Improve MACE-style Finite Model Finding. In *Proceedings of the CADE-19 Workshop: Model Computation - Principles, Algorithms, Applications*, 2003.
- [11] H. Comon. Disunification: A Survey. In *Computational Logic - Essays in Honor of Alan Robinson*, 1991, pp. 322–359.
- [12] T. Eiter, W. Faber, and P. Traxler. Testing Strong Equivalence of Datalog Programs - Implementation and Examples. In C. Baral, G. Greco, N. Leone, and G. Terracina, eds., *Logic Programming and Nonmonotonic Reasoning, 8th International Conference, LPNMR 2005, Diamante, Italy, September 5-8, 2005, Proceedings*, 2005, *LNCs* 3662, pp. 437–441. Springer.
- [13] M. Emmer, Z. Khasidashvili, K. Korovin, and A. Voronkov. Encoding industrial hardware verification problems into effectively propositional logic. In R. Bloem and N. Sharygina, eds., *Proceedings of 10th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2010, Lugano, Switzerland, October 20-23, 2010*, pp. 137–144. IEEE.

- [14] C. G. Fermüller and R. Pichler. Model Representation via Contexts and Implicit Generalizations. In R. Nieuwenhuis, ed., *CADE*, 2005, *LNCS 3632*, pp. 409–423. Springer.
- [15] A. Fietzke and C. Weidenbach. Labelled splitting. In *Ann. Math. Artif. Intell. Vol. 55 No. 1-2*, 2009, pp. 3–34.
- [16] H. Ganzinger and K. Korovin. New Directions in Instantiation-Based Theorem Proving. In *LICS*, 2003, pp. 55–64. IEEE Computer Society.
- [17] G. Higman. Ordering by Divisibility in Abstract Algebras. *Proceedings of the London Mathematical Society*, s3-2(1):326–336, 1952.
- [18] T. Hillenbrand and C. Weidenbach. Superposition for Bounded Domains. In M. P. Bonacina and M. Stickel, eds., *McCune Festschrift*, 2013, *LNCS 7788*, pp. 68–100. Springer. Based on the Research Report MPI-I-2007-RG1-002.
- [19] U. Hustadt, B. Motik, and U. Sattler. Reducing SHIQ-Description Logic to Disjunctive Datalog Programs. In D. Dubois, C. A. Welty, and M. Williams, eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004), Whistler, Canada, June 2-5, 2004*, 2004, pp. 152–162. AAAI Press.
- [20] Z. Khasidashvili, M. Kinanah, and A. Voronkov. Verifying equivalence of memories using a first order logic theorem prover. In *Proceedings of 9th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2009, 15-18 November 2009, Austin, Texas, USA*, 2009, pp. 128–135. IEEE.
- [21] K. Korovin. iProver - An Instantiation-Based Theorem Prover for First-Order Logic (System Description). In Armando et al. [1], pp. 292–298.
- [22] K. Korovin. Inst-Gen - A Modular Approach to Instantiation-Based Automated Reasoning. In Voronkov and Weidenbach [40], pp. 239–270.
- [23] K. Korovin. Non-cyclic Sorts for First-Order Satisfiability. In P. Fontaine, C. Ringeissen, and R. A. Schmidt, eds., *FroCos*, 2013, *LNCS 8152*, pp. 214–228. Springer.
- [24] J.-L. Lassez and K. Marriott. Explicit Representation of Terms Defined by Counter Examples. *J. Autom. Reasoning*, 3(3):301–317, 1987.
- [25] H. R. Lewis. Complexity Results for Classes of Quantificational Formulas. *J. Comput. Syst. Sci.*, 21(3):317–353, 1980.
- [26] W. McCune. Mace4 Reference Manual and Guide. *CoRR*, cs.SC/0310055, 2003.
- [27] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Abstract DPLL and Abstract DPLL Modulo Theories. In F. Baader and A. Voronkov, eds., *Logic for Programming, Artificial Intelligence, and Reasoning, 11th International Conference, LPAR 2004, Montevideo, Uruguay, March 14-18, 2005, Proceedings*, 2004, *LNCS 3452*, pp. 36–50. Springer.

- [28] R. Nieuwenhuis and A. Rubio. Paramodulation-Based Theorem Proving. In Robinson and Voronkov [35], pp. 371–443.
- [29] H. de Nivelle and J. Meng. Geometric Resolution: A Proof Procedure Based on Finite Model Search. In U. Furbach and N. Shankar, eds., *IJCAR*, 2006, *LNCS 4130*, pp. 303–317. Springer.
- [30] J. A. N. Pérez and A. Voronkov. Encodings of Bounded LTL Model Checking in Effectively Propositional Logic. In F. Pfenning, ed., *Automated Deduction - CADE-21, 21st International Conference on Automated Deduction, Bremen, Germany, July 17-20, 2007, Proceedings*, 2007, *LNCS 4603*, pp. 346–361. Springer.
- [31] J. A. N. Pérez and A. Voronkov. Proof Systems for Effectively Propositional Logic. In Armando et al. [1], pp. 426–440.
- [32] J. A. N. Pérez and A. Voronkov. Planning with Effectively Propositional Logic. In Voronkov and Weidenbach [40], pp. 302–316.
- [33] R. Pichler. Explicit versus implicit representations of subsets of the Herbrand universe. *Theor. Comput. Sci.*, 290(1):1021–1056, 2003.
- [34] R. Piskac, L. M. de Moura, and N. Bjørner. Deciding Effectively Propositional Logic Using DPLL and Substitution Sets. *J. Autom. Reasoning*, 44(4):401–424, 2010.
- [35] J. A. Robinson and A. Voronkov, eds. *Handbook of Automated Reasoning (in 2 volumes)*. Elsevier and MIT Press, 2001.
- [36] J. P. M. Silva and K. A. Sakallah. Conflict Analysis in Search Algorithms for Satisfiability. In *ICTAI*, 1996, pp. 467–469.
- [37] J. K. Slaney. FINDER: Finite Domain Enumerator - System Description. In A. Bundy, ed., *CADE*, 1994, *LNCS 814*, pp. 798–801. Springer.
- [38] M. Suda, C. Weidenbach, and P. Wischniewski. On the Saturation of YAGO. In J. Giesl and R. Hähnle, eds., *Automated Reasoning*, Edinburgh, UK, 2010, *Lecture Notes in Artificial Intelligence*, vol. 6173, pp. 441–456. Springer.
- [39] T. Tammet. Finite Model Building: Improvements and Comparisons. In *In: Model Computation Principles, Algorithms, Applications, CADE-19 Workshop W4*, 2003.
- [40] A. Voronkov and C. Weidenbach, eds. *Programming Logics - Essays in Memory of Harald Ganzinger*, 2013, *LNCS 7797*. Springer.
- [41] C. Weidenbach. Combining Superposition, Sorts and Splitting. In Robinson and Voronkov [35], pp. 1965–2013.
- [42] J. Zhang and H. Zhang. SEM: a System for Enumerating Models. In *IJCAI*, 1995, pp. 298–303. Morgan Kaufmann.